




Linux Handbook

A Guide to IBM Linux Solutions and Resources

Understand the Linux implementation
for each IBM @server platform

Review the solutions that IBM
offers for Linux

Consult extensive IBM resources
to implement the Linux solution



Nick Harris
François Armingaud
Massimiliano Belardi
Cameron Hunt
Mauricio Lima
William Malchisky Jr.
Jose Rodriguez Ruibal
Jim Taylor

ibm.com/redbooks

Redbooks



International Technical Support Organization

**Linux Handbook: A guide to IBM
Linux Solutions and Resources**

April 2004

Note: Before using this information and the product it supports, read the information in “Notices” on page xi.

Second Edition (April 2004)

© Copyright International Business Machines Corporation 2003, 2004. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	xi
Trademarks	xii
Preface	xv
The team that wrote this redbook	xv
Become a published author	xxiii
Comments welcome	xxiii
Forward	xxv
Chapter 1. Introduction to Linux	1
1.1 The foundation for Linux	2
1.1.1 Historical perspective	2
1.1.2 UNIX and the culture of collaboration	2
1.1.3 GNU and free software	3
1.1.4 Linux	3
1.2 The IBM commitment to Linux	5
1.2.1 IBM Linux projects	6
1.3 Linux: The operating system	14
1.3.1 Linux BIOS management	16
1.3.2 Where Linux fits in	17
1.3.3 Working with other operating systems	18
1.4 Infrastructure	19
1.5 Data storage	19
1.6 Middleware	20
1.7 Application solutions	22
1.8 Clients	23
1.9 The Linux distributions	24
1.9.1 Inside the distributions	26
1.9.2 What is common among the distributions	28
1.9.3 What is open source	29
1.9.4 The Linux Standard Base project	30
1.10 Red Hat	30
1.11 UnitedLinux	32
1.11.1 What's in the distribution	32
1.12 Conectiva	34
1.13 SuSE	35
1.13.1 The distribution	36
1.14 Turbolinux	37

1.14.1	What's in the distribution	37
1.15	Other players	38
1.15.1	Yellow Dog Linux	41
1.15.2	Red Flag Linux	41
1.15.3	Debian Linux	42
1.16	Going nuts with kernels	43
Chapter 2. Open source software		45
2.1	Creating the code and protecting it	46
2.1.1	Source code	46
2.1.2	Software licenses	48
2.2	The origin of modern open source	49
2.2.1	The Free Software Foundation	50
2.2.2	The compromise of open source software	50
2.2.3	Why the OSS model works so well	53
2.3	Linux, Linux, and open source	54
2.3.1	The GNU utilities	54
2.3.2	The problem of complexity	55
2.3.3	The benevolent dictator	56
2.4	IBM and open source	56
2.4.1	IBM and open source licenses	57
2.4.2	IBM and open source software	57
Chapter 3. Linux, UNIX, and Windows systems		59
3.1	Cost of ownership	60
3.1.1	Total cost of ownership	60
3.1.2	Operating environment and product TCO perspective	62
3.1.3	Findings	63
3.1.4	TCO versus ROI	64
Chapter 4. zSeries and Linux		65
4.1	Introduction: Linux on zSeries	66
4.1.1	Advantages on Linux on zSeries	67
4.1.2	Ways to run Linux on zSeries	68
4.1.3	31-bit and 64-bit options	70
4.1.4	Integrated facility for Linux	70
4.1.5	IPLA pricing	71
4.2	S/390 and zSeries architecture and features	72
4.2.1	Architecture	72
4.2.2	Partitioning and virtualization	74
4.2.3	Networking	76
4.2.4	zSeries FCP support	77
4.2.5	On/Off Capacity	78
4.3	All about z/VM	79

4.3.1	Hardware requirements for z/VM	82
4.3.2	Virtual networking with z/VM	82
4.3.3	VM LAN support	84
4.3.4	z/VM for Linux performance: Scalability	87
4.3.5	Productivity: Development and test workloads	88
4.3.6	Productivity: Production workloads	89
4.3.7	Data backup with IBM Tivoli Storage Manager and FlashCopy	91
4.3.8	System management: Tivoli for Linux on zSeries	93
4.3.9	Consolidation example: Server farm in a box	94
4.3.10	Data and application servers on a single zSeries server	95
4.3.11	Horizontal growth: Adding another server	97
4.3.12	Systems hardware management	98
4.4	zSeries Linux distributions	99
4.5	Linux application solutions	100
4.5.1	Mail serving scenario	105
4.5.2	News group serving scenario	108
4.5.3	IBM Software for Linux on zSeries	111
4.6	How IBM customers are using Linux	117
4.6.1	Customer scenario with SAP on zSeries	118
4.6.2	Local Government, Italy	119
4.6.3	The largest z990 for a Linux solution	121
4.7	zSeries services solutions	121
4.8	Why Linux for zSeries is on demand	124
Chapter 5. xSeries and Linux		127
5.1	X-Architecture™	128
5.1.1	Principles	128
5.1.2	Context	130
5.1.3	Models	130
5.1.4	Linux and functions of X-Architecture	133
5.1.5	Enterprise X-Architecture with Linux	134
5.1.6	xSeries Linux distributions	136
5.1.7	The Red Hat: IBM announcement of December 2003	136
5.2	Intel-based server customers and Linux	137
5.3	Consolidation	138
5.4	Linux clusters	142
5.4.1	High-availability cluster	142
5.4.2	Load-balancing cluster	143
5.4.3	High-performance computing	144
5.5	Linux in a distributed enterprise	146
5.5.1	Basic implementation strategies	147
5.5.2	Business benefits	147
5.5.3	Customer example: A large online company for trading stocks, bonds	

and other personal investment activities	147
5.6 Linux application solutions	149
5.6.1 IBM Software for Linux on xSeries	149
5.6.2 ISVs and application solutions	150
5.6.3 Customer scenario: A leading company in Web-based time- and expense- tracking solutions.	158
5.6.4 Customer scenario: LOSSRUNS Inc.	160
5.7 Infrastructure solutions	161
5.7.1 Customer scenario: Butterfly.net.	165
5.8 xSeries architected solutions.	166
5.8.1 Blade server	166
5.8.2 Customer scenario: Wolferman's	167
5.9 xSeries services solutions.	170
5.9.1 Implementing Linux.	170
5.9.2 Support	171
Chapter 6. pSeries and Linux	173
6.1 How Linux runs on the pSeries server	174
6.1.1 Linux-capable pSeries model structure.	174
6.1.2 pSeries operating system scenarios	176
6.1.3 AIX 5L and Linux.	177
6.1.4 Options to run Linux on the pSeries server.	180
6.1.5 Linux and its relationship with AIX	184
6.2 Linux for pSeries hardware enablement	184
6.2.1 Linux for pSeries scalability.	185
6.2.2 Linux for pSeries's Reliability, Availability, and Serviceability.	186
6.2.3 Other sources of information.	187
6.3 Consolidation.	190
6.4 Linux for pSeries clusters	191
6.5 Linux for pSeries application solutions	192
6.5.1 IBM Software for Linux on pSeries	192
6.5.2 Other Linux for pSeries software applications.	194
Chapter 7. iSeries and Linux	197
7.1 How Linux runs on the iSeries server	198
7.1.1 Why Linux on iSeries	199
7.1.2 Linux and LPAR	201
7.1.3 iSeries Linux integration	204
7.1.4 Linux on iSeries models	204
7.1.5 iSeries Linux distributions	209
7.2 Consolidation.	211
7.3 Linux clusters	212
7.4 Performance	213

7.5 Linux in a distributed enterprise	214
7.6 IBM Software Group products for iSeries	215
7.7 iSeries ISV application solutions	218
7.8 iSeries solutions	230
7.9 iSeries service solutions	234
7.9.1 iSeries Linux resources	234
Chapter 8. TotalStorage and Linux	237
8.1 Linux and IBM TotalStorage products	238
8.1.1 Why storage solutions	238
8.2 IBM TotalStorage solutions	241
8.3 IBM TotalStorage product line	242
8.3.1 IBM TotalStorage disk products	242
8.3.2 IBM TotalStorage tape products	244
8.3.3 Comparing SAN X NAS technologies	245
8.3.4 IBM TotalStorage NAS products	246
8.3.5 IBM TotalStorage SAN products	250
8.3.6 IBM TotalStorage Proven™ Program	253
8.3.7 ESS	253
8.3.8 FASTT	261
Chapter 9. IBM Software for Linux	265
9.1 Overview of Linux and IBM Software availability	266
9.2 On Demand business	268
9.2.1 Open standards	269
9.2.2 Integration	270
9.2.3 Automation	270
9.2.4 Virtualization	271
9.3 IBM WebSphere brand products on Linux	271
9.3.1 WebSphere product availability on Linux	274
9.3.2 WebSphere Portal for Multiplatforms	274
9.3.3 Linux-related WebSphere resources	276
9.4 IBM Information Management products on Linux	279
9.4.1 DB2 product family	279
9.4.2 Informix product family	282
9.4.3 U2 product family	283
9.4.4 IBM Information Management resources	284
9.5 IBM Tivoli products on Linux	286
9.5.1 Tivoli product availability on Linux	290
9.5.2 Tivoli resources	291
9.6 Lotus products on Linux	292
9.6.1 Lotus product availability on Linux	292
9.6.2 Lotus resources	294

9.7 Rational Software products on Linux	295
9.7.1 Rational product availability on Linux	296
9.7.2 Rational resources	297
9.8 Additional IBM Software	298
9.8.1 Product availability on Linux	298
9.8.2 Additional resources	298
9.9 IBM Integrated Platform for e-business	299
9.10 IBM Software on Linux deployment examples	301
9.10.1 Using WebSphere Application Server on the zSeries server	301
9.10.2 Using Domino and WebSphere on xSeries and iSeries	302
9.10.3 DB2 UDB, DB2 Connect, and WebSphere Portal	303
9.10.4 Business continuity solution for Disaster Recovery	304
9.11 IBM and Linux solution customer success	305
9.12 Additional IBM Software for Linux resources	306
9.12.1 Evaluation and demonstration software	307
9.12.2 IBM Software on Linux training	308
9.12.3 IBM Redbooks	308
9.12.4 IBM Software support organizations	309
Chapter 10. Linux clusters	313
10.1 An introduction to clusters	314
10.1.1 Core cluster concepts	314
10.1.2 IBM cluster offerings	323
10.2 Linux cluster components	324
10.2.1 Nodes	324
10.2.2 Networks	327
10.2.3 Support software	331
10.2.4 Applications	332
10.3 Linux cluster history	334
10.4 Linux cluster solutions	338
10.4.1 Cluster 1350	338
10.4.2 Cluster 1600	341
10.4.3 ClusterProven®	342
10.4.4 Solution examples	342
10.5 The future of Linux clusters	346
10.5.1 Emerging trends	346
10.5.2 The role of Linux clusters in emerging trends	349
Chapter 11. IBM services for Linux	351
11.1 IBM Global Services for Linux	352
11.1.1 Assess, plan, and design	352
11.1.2 Implement and run	353
11.1.3 Support: IBM Support Line	356

11.2 IBM Products and Solutions Support Center	359
11.3 IBM Software Services for Linux	365
11.3.1 WebSphere	366
11.3.2 DB2 for Linux	367
11.3.3 Tivoli	368
11.3.4 Lotus	369
11.4 IBM platform services for Linux	369
11.4.1 zSeries services	370
11.4.2 iSeries services	370
11.4.3 pSeries services	371
11.4.4 xSeries services	371
11.4.5 Consolidated platform matrix	372
11.5 Linux assessment services	372
11.5.1 Business challenge	372
11.5.2 Solution	372
11.5.3 At a glance: Solution components	372
11.5.4 Benefits	373
11.5.5 Technical details and architecture	373
11.6 Server consolidation study	373
11.6.1 Highlights	373
11.6.2 Pricing	374
11.6.3 IBM Linux Implementation Services	374
11.7 Examples of Linux services	374
11.7.1 Example 1: Billion-dollar financial institution	375
11.7.2 Example 2: Multi-million dollar advertising company	375
11.7.3 Example 3: One of the largest network services providers	375
Chapter 12. Linux solutions	377
12.1 Total IBM Linux solutions abstracts	378
12.1.1 IBM Global Solutions Directory	378
12.1.2 Computer-aided engineering	378
12.1.3 Linux in exploration and production	380
12.1.4 IBM Lotus Domino Web Access for Linux	381
12.1.5 Real-time securities and risk analysis on Linux	383
Related publications	385
IBM Redbooks	385
Other publications	386
Referenced Web sites	387
How to get IBM Redbooks	387
Help from IBM	387
Index	389

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation and the Rational Software Corporation, in the United States, other countries, or both:

Redbooks (logo)  ™	FlashCopy®	POWER3™
IBM @server®	FICON®	POWER4™
alphaWorks®	GDPS®	QuickPlace®
developerWorks®	HiperSockets™	Rational Rose®
IBM®server®	Hummingbird®	Rational®
e-business on demand™	HACMP™	Red Brick™
e-business Hosting™	Informix®	Redbooks™
ibm.com®	Intelligent Miner™	RedBack®
iNotes™	IntelliStation®	RACF®
iSeries™	IBM TotalStorage Proven™	RS/6000®
pSeries®	IBM®	S/360™
xSeries®	IMS™	S/370™
z/Architecture™	LearningSpace®	S/390®
z/OS®	Lotus Enterprise Integrator®	Sametime®
z/VM®	Lotus Notes®	Seascape®
zSeries®	Lotus Workflow™	SecureWay®
AFS®	Lotus®	Sequent®
AIX 5L™	LANDP®	ServerProven®
AIX®	Monday™	ServeRAID™
APL2®	Multiprise®	SmartSuite®
AS/400®	MQSeries®	SmoothStart™
BladeCenter™	MVS™	Storage Tank™
Blue Gene®	Netfinity®	System/36™
BookManager®	NetServer™	System/370™
Chipkill™	NetView®	System/390®
ClearCase MultiSite®	NetVista™	SystemBuilder™
ClearCase®	Notes®	SANergy®
ClearQuest®	NUMA-Q®	Tivoli Enterprise™
ClusterProven®	OS/2®	Tivoli Enterprise Console®
CICS®	OS/390®	Tivoli®
DataBlade™	OS/400®	TotalStorage Proven™
Domino®	Parallel Sysplex®	TotalStorage®
DB2 Connect™	PartnerWorld®	ViaVoice®
DB2 Universal Database™	PowerPC Reference Platform®	VisualAge®
DB2®	PowerPC 603™	VM/ESA®
DFS™	PowerPC 604™	VSE/ESA™
DFSMS/MVS®	PowerPC®	Wave®
Enterprise Storage Server®	PurifyPlus™	WebSphere®
Everyplace®	PAL®	X-Architecture™
ESCON®	POWER™	1-2-3®

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Preface

IBM® has a strong commitment to Linux as an operating system as well as to Linux-based solutions. This IBM Redbook gives a broad understanding of IBM products and how they relate to the Linux operating system. It tries to encapsulate the IBM Software Solutions that are available for Linux-based servers. It also describes the Linux enablement throughout the IBM @server product line.

This second edition includes a new chapters on Linux Solutions and IBM Storage in a Linux environment. All the other chapters have been revised to cover new IBM announcements, and the amazing growth in Linux adoption.

This IBM Redbook offers enough information to give you a high-level understanding of Linux. Therefore, it is not meant to be a highly technical document. It also provides pointers to further sources of information.

This redbook is intended for IBM representatives, business partners, and customers who are involved in considering and planning Linux-based software solutions. You will find this redbook particularly helpful if you are planning to test or implement Linux in an enterprise.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Rochester Center.



Nick Harris is a Consulting IT Specialist for the iSeries and has pent the last six years in the International Technical Support Organization, Rochester Center. He specializes in LPAR, iSeries hardware and software, external disk, Integrated xSeries® Server for iSeries, and Linux. He writes and teaches IBM classes worldwide on areas of i5, iSeries and AS/400® system design and server consolidation. He spent 13 years in the United Kingdom (UK) AS/400 Business, and has experience in S/36, S/38, AS/400, and iSeries servers. You can contact him by sending e-mail to: <mailto:niharris@us.ibm.com>



Francois Armingaud (FDA) has been with IBM for over 25 years. His last position in IBM was as a member of the EMEA West Security Team. He is concurrently on partial assignment with IBM, while teaching university courses in France. He also continues his consulting projects and work on security and self-healing systems patents. He became involved with Linux in 1995 and has since become a Linux evangelist in IBM France. Francois participated in four previous IBM Redbook projects. His last two projects involved C and C++ programming with Pthreads, in Austin, Texas, and AIX® 4.3 elements of security, in Poughkeepsie, New York. Francois' e-mail address is: <mailto:fda@armingaud.com>



Massimiliano Belardi is an IT Specialist for Linux, working for IBM Global Services since 1998. He has designed system solutions architected on Linux on the IBM @server zSeries® server. He has designed these solutions for many customers in Italy and maintained an active role in these projects to ensure their successful deployment. For assistance with zSeries projects in Italy or in the rest of Europe, you can contact Max by sending e-mail to: <mailto:m.belardi@it.ibm.com>



Cameron Hunt is an IT Consultant and Educator in the United States. He has 10 years of experience in the IT field and has spent the last three years focusing on Linux. His areas of expertise include Linux, Microsoft, and Novell system administration. He has written extensively about computer security.




Mauricio Lima is an Infrastructure Specialist at EDS Brazil. He has five years of experience in the Linux field and nine years of experience in the IT field. He has spent the last three years working in Information Technology Outsourcing at EDS Brazil. For EDS customers, he has designed system solutions architected on Linux and multiplatform environments for many customers in Brazil and Latin America. He has also been actively involved to ensure their successful deployment. Mauricio is certified at LPI Level 1 and Citrix Metaframe for UNIX. His e-mail address is: mail to:mauricio.lima@eds.com



William Malchisky Jr. began Effective Software Solutions, LLC, in 1999, after five years of technology consulting on Wall Street. He expanded his knowledge to Linux after six years of AIX and UNIX development in college. His firm targets Lotus® Notes® projects for Fortune 100 companies. It also performs broad-based business and technology consulting for small businesses in the New England and New York Metro area. His firm uses Linux exclusively for all internal operations. You can contact him by sending e-mail to:

<mailto:william.malchisky@effectivesoftware.com>




Jose Rodriguez Ruibal Jose Rodriguez Ruibal is a Linux Specialist in the Advanced Technical Support (ATS) Products and Solutions Support Center (PSSC) for EMEA in Montpellier, France. Prior to joining the ATS team, he worked for Red Hat in Spain and Portugal and for the National French Institute for Research in Computing and Robotics (INRIA) in Linux-related activities. He has five years of experience in Linux and UNIX-like network operating systems. He holds a bachelor's degree with honors in Computer Science from the Universidad Antonio de Nebrija in Madrid, Spain, and

he is also Red Hat Certified Engineer. His areas of expertise include xSeries hardware and TotalStorage® solutions, and Linux. He has also authored a white paper about Oracle 9i RAC installation under Linux on xSeries x360 and x440 machines, another on Installing Linux on the IBM @server™ xSeries 450, and the redbook *IBM @server xSeries 450 Planning and Installation Guide*. His e-mail is: mailto:pssc_linux@fr.ibm.com



Jim Taylor is an Independent Consultant specializing in AIX, IBM @server pSeries® clusters, and Linux. His experience includes over 30 years in the IT industry and over 14 years in AIX and other UNIX-like operating systems. He retired from IBM Canada in 2000. Currently, Jim teaches IBM courses in AIX, Parallel System Support Programs (PSSP), and Linux, and promotes Linux everywhere. He does consulting work and is a technical editor of books about Linux, UNIX, and Windows. His e-mail address is: <mailto:jttac@sympatico.ca>

We also thank the residents who worked on the first edition of this handbook:



Randy Kuseske is currently assigned to Server Group Market Intelligence at IBM Rochester. He has 21 years of experience in the IT industry, previously working at Unisys, EMC, and AiC. His areas of focus have included operating system internals, compiler and language technology, globalization, and product competitive analysis. He currently performs product competitive analysis on various vendor environments, products, and technologies for iSeries, IBM @server, and IBM Market Intelligence initiatives. His e-mail address is: <mailto:random1@us.ibm.com>



Dan Lacine is an Advisory Education Specialist at the IBM Rochester Support Center. He provided remote support for 21 years ranging from System/390® software to iSeries Communications and Client Access. He currently develops and teaches classes to IBM personnel on iSeries products such as iSeries Navigator and iSeries Access for Web, and PC-based operating systems. He co-authored a redbook titled *Make Your AS/400 Year 2000 Ready*, SG24-5401, and has written articles for *AS/400 Magazine* and the *UK Assist/400* technical newsletter.




Tomomi Takada is an IT Specialist at the Linux Support Center of IBM Japan. He has three years of experience in the Linux field. He has provided technical support for the public sector and been in charge of IBM @server xSeries servers. His areas of expertise include Linux, xSeries, security, and systems management for Intel-based servers. His e-mail address is: <mailto:e30810@jp.ibm.com>



Dirk Webbeler is an IT Specialist in Germany. He has five years of experience in iSeries field and more than 10 years in the IT business. Working in sales and technical support, his areas of expertise include TCP/IP, local area network (LAN) and wide area network (WAN) planning, Windows NT and 2000, server consolidation, implementation of Management Information Systems (MIS), and database design. He has written extensively on MIS, dynamic programming, and Windows NT/2000 Server consolidation. His e-mail address is:

<mailto:Dirk.Webbeler@dw-solutions.de>



Alexander Zaretsky is an IT Support Specialist at the IBM Andean pSeries Software Support Center in Bogotá, Colombia. He has more than ten years of experience in the UNIX system and C programming fields, and five years as an AIX and Linux System Administrator. His areas of expertise include programming and development in AIX and Linux environments, Tivoli® product support, and network configuration and security. He holds a degree in computer systems engineering from The Colombian School of Engineering. His e-mail address is:

<mailto:alexzare@co.ibm.com>



Ralph Cooley is a Linux Sales Specialist at IBM and has been an IBM employee since the purchase of Sequent® Computer Systems in 1999 by IBM. He promotes Linux and the IBM Linux initiative to the public and to IBM prospects, customers, and partners. Ralph has been a Linux user for four years. He has 25 years of UNIX market experience as a system administrator, application developer, kernel developer, instructor, consultant, and salesperson. Ralph holds a Bachelor of Science in ocean engineering and has earned a master's in Business Administration. His e-mail address is: <mailto:cooleyr@us.ibm.com>

Bob Waite is an Advisory Software Engineer. He has worked at IBM Rochester for 24 years writing product documentation for User Technologies. Currently, he is the documentation team leader for Linux and Java on iSeries. Previously, he was the documentation team leader for WebSphere®, San Francisco, and other object technologies. His e-mail address is: <mailto:rwaite@us.ibm.com>

Thanks to the following people for their contributions to this project:

Scott Knupp
Worldwide Program Director, Linux Sales

Whei-jen Chen, Data Management
Rufus Credle, xSeries and WebSphere
Maritza M. Dubec, Technical Editor
Greg Geiselhart, zSeries
Pete Kovari, WebSphere
Edson Manoel, Clusters
David Watts, xSeries
IBM ITSO

Dave Boutcher
Michael P Martin
Craig Johnson
iSeries

Jay Ashford
Chuck Bryan
Ron Gordon
pSeries

Carol Carson
Amy Freeman
Michael T Keleher
Cindy Roy
xSeries

Ed Gauthier
Susan Greenlee
zSeries

Joe Banas
Doug McGuire
Richard C Ferri
Egan Ford
John Hawkins
Norm Snyder
Mark Solomon
Wayne Wellik
Clusters/Grid

Cristina Zabeu
Mark Flemming
IBM TotalStorage

Peyen Fong
Daren Hanson
Adam Jollans
IBM Software Group

Judith Adolino
Tony Ravinsky
IBM Global Services

Bob Butler
Ron Houston
Yohichiroh Ishizak
Dr. Harald Neumann
Joseph Rhoden
IBM Linux Distribution Partner Alliances

Boris Barth
mySAP Technologies on zSeries Planning & Marketing

Claudia Prawirakusumah
Ilse Schmidt
IBM@server Market Planning

Gerald Hosch
WW Marketing Linux on zSeries

Ken Godfrey
Linux Solutions

Pete Hnath
Red Hat

Takashi Nishioka
Turbolinux

Martina Krahmer
SuSE

We would like to thank Alfred Sasser of Xerox for his work on the Penguins cover on the special edition, and the rest of the Xerox printing team for their hard work producing the hardcopies.

Many people sent in reviews of this book in its draft stages and first edition, plus they provided feedback through IBM Redbooks™. With this second edition, we look forward to more valuable input from you in the future.

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- Send your comments in an Internet note to:

redbook@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. JLU Building 107-2
3605 Highway 52N
Rochester, Minnesota 55901-7829

Forward

In this handbook, we set out to describe in tangible terms, the breadth and depth of IBM's commitment to Linux as well as the broad spectrum of choice offered by IBM in terms of alternative Linux architectures, platforms, and solutions. At the same time, we wanted to provide a solid background for business and technical leaders seriously considering Linux for implementing a new class of service oriented enterprise solutions.

IBM's early entry into the Linux marketplace was due to an expectation that Linux would provide unprecedented choice, value, and flexibility for our customers and partners. This initial expectation has played out perfectly. Today, IBM's broad investment in Linux across our business is reflected in the vast number of IBM product, services, and support offerings available in the marketplace. Satisfied customers further validate our strategy! Linux continues to be the fastest growing server operating system in the industry.

The Linux ecosystem is large, diverse, and increasing rapidly. IT surveys indicate that a rapid adoption of Linux based systems and technologies is underway. Around the world, governments are choosing Linux for independence and security. Independent Software Vendors (ISVs) have embraced Linux as a viable platform for growing their business'. The Open Source Development Labs (OSDL) has just released a much anticipated 2.6.0 Linux kernel. IBM continues to regularly announce new products and offerings relevant to Linux such as: Blade Servers, 64 bit computing with the POWER™ architecture, Linux client services, and many more discussed within this handbook.

The chapters which follow are a result of the sustained investments that IBM is making in Linux through the open source development community, through IBM's own world-wide development labs, services organizations, and partner support programs. IBM continues to lead the industry by delivering world class Linux products, services, and support on a global basis!

I sincerely hope that this second version of the Linux handbook provides you with both the knowledge and confidence to implement your next project with an IBM Linux solution!

Dr. D. Scott Knupp

Worldwide Program Director
Linux Sales



Introduction to Linux

This chapter examines the components and features of Linux, and its application and middleware support. It reviews Linux capabilities in both the server and client environments. Plus, it looks at what comprises a distribution, and highlights the major features of the most popular distribution.

1.1 The foundation for Linux

This chapter follows the history of the UNIX operating systems to help you better understand how Linux came about. It also provides a more concrete overview of the beginning of this revolutionary operating system.

1.1.1 Historical perspective

Years ago when the time came to install an operating system, a good plan was to have some activity to occupy your time during the tedious phases. With the advent of fast processors, high-bandwidth networks, and high-speed drives for installation media, the luxury of “popping in a tape and going to lunch” was reduced to “kickstart the installation and read the splash screens for ten minutes.” You could choose to save this chapter as a possible diversion during these ever-shorter installation periods.

Nowadays, nearly everything you read about Linux begins with the ritual invocation of “... an operating system created in 1991 as a hobby by Linus Torvalds at the University of Helsinki...” While this is true, it does not do justice to the significance of the work and its broad implications for today’s implementers. We know that Linux refers to a UNIX-like operating system. Therefore, we begin with a brief overview of the development of portable operating systems and open source to see the context in which Linux evolved.

1.1.2 UNIX and the culture of collaboration

In 1969, several AT&T Bell Labs employees began work on an operating system that would come to be called *UNIX*. A significant novelty of their development was that the operating system was portable across hardware platforms. Prior to this, it was more typical to emulate the (usually) older hardware on a new system, so the operating system would run unchanged, or to rewrite the operating system completely for the alternate hardware. Such portability was achievable only through writing most of the operating system in a higher level language “above” the hardware. Unlike an operating system written in Assembly language for a particular architecture, this abstraction and the language (C) which they developed to implement it permitted the study of their operating system without much regard to hardware specifics. And by 1976, UNIX was being taught in classes on operating systems at the university level.

At the time, AT&T for various legal reasons permitted free academic access to the source code to UNIX while charging over \$20,000 U.S. (in 1976 dollars) for commercial or government access. AT&T halted publication of the source code in university texts since this revealed proprietary Bell Labs code. The era of collaborative programming had arrived.

1.1.3 GNU and free software

In the U.S., the academic variant of interest became the Berkeley Systems Distribution (BSD), where virtual memory and networking were added. These advancements permitted large collaborative projects with contributors being scattered throughout the world. Lawsuits eventually ensued among AT&T, the Regents of the University of California, and other parties over access to and distribution of the operating system source code. Such constraints on intellectual property rights to code, along with commercialization of academic artificial intelligence projects in the late 1970s and early 1980s, provided strong motivation for one researcher from the Artificial Intelligence Laboratories at the Massachusetts Institute of Technology to write an operating system that was portable. This system would also be licensed in a manner to prevent its eventual constraint by intellectual property claims.

The new operating system was to be named GNU, a recursive acronym for “Gnu’s Not UNIX.” It would be licensed under the GNU General Public License (GPL). This was the birth of “free (as in freedom) software”, in contrast to software in the public domain. This is free as in “freedom of speech” not free as in “free food.” To understand the “Four Freedoms” and rules of the GNU GPL, visit the GNU Web site at:

<http://www.gnu.org/philosophy/free-sw.html>

By this time, such vendors as Sun, Hewlett-Packard, and IBM had proprietary commercial offerings derived from licensed AT&T UNIX. These offerings were gaining popularity with corporate customers. The GNU development effort began by making tools such as editors, compilers, and file utilities in source form that could be compiled and executed on any platform. They standardized and improved upon those offered by commercial vendors.

Around 1990, programmers had contributed a nearly complete operating environment to GNU with the exception of a kernel. The GNU kernel was to be based on a microkernel architecture for improved portability. This approach required an (arguably better) architecture that was completely different from that of a monolithic kernel. The GNU kernel project is known as the *Hurd*.

1.1.4 Linux

With the withdrawal of AT&T source code from the university environment, and the BSD operating system mired in legal challenges by AT&T to its claim to be unencumbered by proprietary code, a small scale UNIX-like skeleton of an operating system called *Minix* was published in a text to be used as a teaching tool.

It is here that Linus Torvalds enters the story. He decided to write a UNIX-like operating system with improved functionality over that of Minix to run on readily available personal computers. His purpose was to teach himself the C language and improve his understanding of operating system design. He and colleague Lars Wirzenius published their source code under the GPL on the Internet for public comment, and Linux was born.

Linux was a kernel without utilities. GNU was an operating environment lacking a finished kernel. And unencumbered non-kernel BSD pieces were available to complete the picture. In short order, the components were combined with installation and maintenance tools and made available by distributors. The first of serious note was Slackware in 1993, which followed by many others, making the GNU/Linux (or simply Linux, as it has come to be known) combination readily available. In only a few years, a worldwide Linux community evolved, comprised of programmers and users attracted by the reliability and flexibility of this “free” operating system.

The term *open source* began to replace the term *free software* as the commercial adoption of GNU/Linux grew. There is in fact a difference upon which hinges the fate of commercial ventures in this arena.

Maybe one of the most interesting aspects of the beginning of Linux is that Torvalds invited programmers from around the world to comment on and to improve his code. This is one of the key ideas behind the success of Linux. With the world as the laboratory, the number of testers and developers is nearly endless. It is because of this resource that Linux is constantly evolving and improving.

Since the Linux source code is freely available, several companies have developed different *distributions* of Linux. A distribution is a complete system. The key component is the Linux *kernel*. Other utilities, services, and various applications can be included as well, depending on the distribution and the intended use. There is no standard distribution. Each distribution that is available has unique advantages.

IBM was early to recognize the value of Linux:

- ▶ Investing in Linux-related product development
- ▶ Forming alliances with key Linux distributors
- ▶ Contributing to the open-source community
- ▶ Aggressively supporting the platform

IBM believes this investment will benefit its customers as they continue to exploit Linux for their IT infrastructures and e-business.

Figure 1-1 shows the Linux ecosystem chart in which IBM is a key component of life in the Linux world.

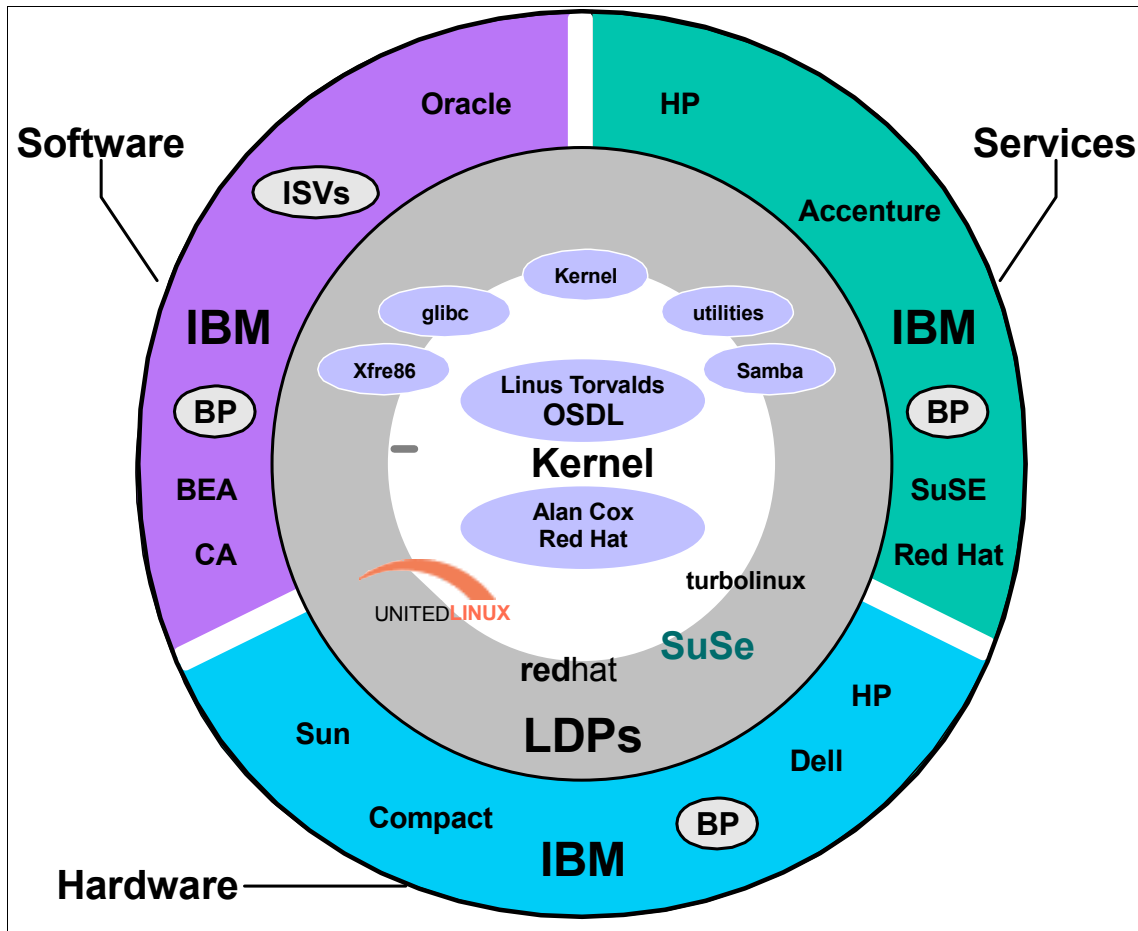


Figure 1-1 The Linux ecosystem

1.2 The IBM commitment to Linux

IBM is fully committed to the open source movement and believes that Linux is a key platform for e-business. IBM works with the open-source community. We leverage relevant technologies and experience to help enhance Linux, to define the standards, and to extend Linux to the enterprise level.

IBM provides continued support and participation throughout the world. You can learn more about this support and IBM commitment to Linux at the following Web sites:

- ▶ The IBM Linux portal for a general point of entry into IBM and Linux
<http://www.ibm.com/linux>
- ▶ IBM Linux Technology Center (LTC)
<http://www.ibm.com/linux/ltc>
- ▶ IBM Solution Partnership Centers
<http://www.developer.ibm.com/spc/index.html>
- ▶ IBM Linux support line
http://www.ibm.com/services/e-business/linux_8.html
- ▶ The Open Source Development Lab
<http://www.osdl.org/>
- ▶ IBM developerWorks® Linux
<http://www-106.ibm.com/developerworks/linux/>
- ▶ IBM alphaWorks®
<http://www.alphaworks.ibm.com/>

As part of this continuing commitment, IBM has teamed with leading commercial Linux distributors including Red Hat, SuSE, and Turbolinux. This cooperation allows IBM to port, test, and certify the performance of its offerings running on various Linux distributions. The result is to enable you to exploit the full potential of Linux.

1.2.1 IBM Linux projects

At IBM, there is a real thrust to adopt Linux technology where it makes sense. The following sections highlight some internal IBM projects that use 2000 Linux servers and tens of thousands of Linux clients in our corporate environment.

Web infrastructure

Development of IBM Web infrastructure is expanding with the use of Linux in the following projects:

- ▶ Linux portal
- ▶ Web content management system
- ▶ Advanced search engine
- ▶ W3 e-workplace development environment
- ▶ Intranet forums

- ▶ W3 e-workplace special events
- ▶ Hone/IBM link Web portal

Linux portal

The external IBM Linux portal deployed on two redundant, load balanced xSeries servers running Apache Server on Red Hat Linux with these four additional products:

- ▶ IBM HTTP Server
- ▶ IBM Tivoli Storage Manager Client and IBM Tivoli Monitoring Agent (TMA)
- ▶ Server Resource Management (SRM)
- ▶ Enterprise Security Manager (ESM)

The Linux portal is located at:

<http://www.ibm.com/linux>

Linux Solutions portal

The external IBM Linux Solutions portal is an extension of the Linux portal that offers an overview of the Linux solutions supported and offered by IBM. This portal is a direct link to find information about Linux solutions that use IBM hardware or software, supported by IBM services specialists or Business Partners.

The Linux Solutions portal is located at:

<http://www.ibm.com/Linux/solutions>

Web content management system

The IBM Linux portal team developed a content management system. This system uses Linux to facilitate the development, delivery, and tracking of the content for the <http://www.ibm.com/linux> site. This includes a backend DB2® server running on Linux on zSeries hardware.

Advanced search engine

The W3 Advanced Search facility has been upgraded to use an IBM developed search engine, which is derived in part from our WebFountain and Juru research. The search engine is now hosted on five SMP xSeries servers running Linux as well as TSM Client, Tivoli Monitoring Agent (TMA), Server Resource Management (SRM), Enterprise Security Manager (ESM), WebSphere Application Server, and DB2.

W3's On Demand workplace development environment

The IBM intranet application (W3) development team is developing all of their applications on development and test environments on zSeries virtual machine (VM) guests running Linux.

Intranet forums

This project replaced the VM Forum's infrastructure with two Linux logical partitions (LPARs) on each of the two zSeries servers that are accessible from the IBM W3 intranet. IBM used VM Forums for over 20 years.

This collaboration system provides service to over 350,000 IBM representatives worldwide. It is now a repository of news groups with over 1,000 topics that have around 15,000 new posts every month.

W3 On Demand special events

Linux hosted W3 *jam sessions*. *Manager Jam* was hosted on zSeries hardware, and ran the event on five VM virtual servers running SuSE Linux. These servers ran IBM HTTP Server and WebSphere Application Server 3.5, and one of the images ran DB2 and IBM HTTP Server.

Jamming is the art of strategic brainstorming across a corporation. It is a technique IBM Research began using about a year ago. The term originates from the jazz and blues music concept of a jam session, which is the ultimate gathering of expert musicians to create something great. No one knows what will be produced during a jam session, only that it will be something that has never been heard before.

In a way, this is what IBM jam sessions are about. We gather experts, give them a scenario to ponder or a problem to solve, and then set them loose. Those experts, like a group of great musicians, then create a series of "riffs" or great ideas that may not surface through more traditional problem-solving processes.

Hone/IBM Link Web Portal

The portal Web site for EMEA's customer order support functions was implemented on a set of 15 Linux guest images under VM on a s/390 machine running Linux S/390® WebSphere server. This configuration includes Enterprise Access Enabler to VM legacy applications and databases. The functions include inventory retrieval, order submission, price inquiry, WW publications, and internal access through HoneWeb. After the successful pilot phase, this project is in development for full production.

Security

Security is always an important project area. The following security projects are Linux based:

- ▶ Security assessments
- ▶ Virus detection
- ▶ Storage Architecture Security Directory
- ▶ E-mail antivirus and antispyware scanners

Security assessments

IBM Global Services Internet Vulnerability Scanning Services are handled by 112 Intel boxes running Linux. Through this configuration, IBM Global Services provides internal and commercial Internet and intranet vulnerability scanning services. IBM serves over 100 commercial customers and scans 30,000 internal IP addresses per week through this service. Linux has used these services for over five years.

Virus detection

Linux servers have been running “honey pot” and scanning software for security purposes to detect and track attempted attacks and viruses. One server is up and running in each of the 13 Asia Pacific countries (or regions).

Honey pot: This computer system on the Internet is expressly set up to attract and “trap” people who attempt to penetrate other people’s computer systems.

Storage Architecture Security Directory

Our internal Global Storage Architecture calls for a Linux based security directory that contains user IDs, groups, and passwords for all its users. There are also Linux based performance monitoring servers deployed in this infrastructure.

E-mail antivirus and antispam scanners

This project deployed antivirus and antispam scanning software into the IBM mail infrastructure. This project involves scanning both incoming and outgoing mail for viruses and spam. Trend Micro Interscan VirusWall is the antivirus software that was selected. They are running on 44 Linux xSeries machines worldwide.

Monitoring

The system monitoring projects include:

- ▶ Network monitoring
- ▶ Performance monitoring
- ▶ Asset monitoring
- ▶ Operations for e-hosting and network management

Network monitoring

IBM Global Services, Networking Services, e-business hosting services has consolidated their Netview network monitoring servers from 62 separate NT and AIX servers onto a shared z800 with two IFL engines shared between two LPARs using virtual Linux guests under zVM. This has reduced the hosting costs by 40% with room for growth that could lead to a 60% savings.

Performance monitoring

Linux supports the IBM End-to-End Probe Platform (EPP), a Java-based monitoring tool primarily built to run on the Windows NT platform, but that also runs on Linux. EPP measures the performance and availability of applications from an end user's perspective. The performance and availability metrics are generated by periodically executing and measuring the duration of emulated transactions. Each probe server instance can poll approximately 180 servers per hour using the default time-out values.

There are 19 EPP servers running Linux, deployed in 13 different countries in EMEA. There are six in the U.S.A. and four in Canada. Linux is an ideal operating system for this kind of application because it is reliable and economical. Also, its multitasking capabilities provide four times the throughput over the Windows NT solution. This allows IBM to use 75% fewer servers for the same workload.

Asset monitoring

The workstation asset management tools Asset Center, Brio Reporting, and Communication Director track and manage workstation hardware and software assets. These tools run on a Linux image that is running as a guest under z/VM® on S/390 hardware. There is also a DB2 database on the S/390 Linux for software inventory.

Operations for e-hosting and network management

Linux is used to monitor over thousands of commercial URLs hosted by IBM Global Services. It is also used for remote console and cluster management. IBM uses Linux to monitor thousands of network interfaces.

File and print

The two file and print projects are:

- ▶ File serving
- ▶ File and print servers

File serving

IBM Global Services is using a network of Linux file servers to provide standard client images worldwide through IBM Standard Client Installer (ISCI). It also provides the IBM Standard Software Installer (ISSI) application and maintenance fix packages to end users. There are currently 168 ISCI/ISSI servers in production in 32 countries (or regions) worldwide.

File and print servers

Australia and New Zealand converted 66 file and print servers previously on OS/2®, Windows 95, and Windows NT to 44 xSeries servers running Linux and one RS/6000® box running AIX 5.1L across 13 sites.

Follow-on projects are probable in other Asia Pacific countries (or regions). This project included a Linux domain controller for each site as well.

IBM's Global Voice Infrastructure

IBM is using Linux in a SIP proxy server to support using Voice over IP technology. It is also using Linux in a pilot softswitch solution for a planned 2004 rollout.

IBM manufacturing and development facilities

Manufacturing and development facilities include six projects:

- ▶ Microelectronics 300mm wafer manufacturing and test
- ▶ Electronic design automation
- ▶ VLSI test engineering
- ▶ Manufacturing line kiosks
- ▶ Software development
- ▶ IBM Internal Open Source Bazar

Each of these projects is discussed in the following sections.

Microelectronics 300mm wafer manufacturing and test

The future 300mm wafer chip manufacturing line tool controllers from the IBM Microelectronics Division run on xSeries machines running Linux. They were slated to be deployed with Windows 2000, but a major reliability concern was detected.

This project demonstrated that Linux was much more reliable in this environment and showed a 10% performance gain. The chip manufacturing line tool controllers currently run on 200 xSeries boxes running Linux. Two hundred to 300 installations are planned across both the East Fishkill and Burlington, New York, sites. After almost a full year of production support, there have been no Linux-related outages.

Electronic design automation

IBM Microelectronics Division embarked on a multiyear strategy to add Linux as a key element of their design automation platform. This involves a long-term migration of design workstations and development infrastructure to Linux, as well as tools and applications porting. There are 10 servers now in production.

To date, already 80% of the tools have been ported, with a year-end completion goal. This division is also developing a developer/designer workstation environment. This is based on a vision for a potential for productivity gains, performance benefits, and functional benefits of Linux-based tools.

VLSI test engineering

The microelectronics chip manufacturing and test engineering organization ported their test data generation programs to Linux. This enabled the use of two xSeries servers in the process of generating device test patterns and development of programs for the chip testers in the manufacturing process.

Manufacturing line kiosks

IBM Canada's Bromont manufacturing facility installed 81 kiosks to access Web applications, legacy host applications, and printers using old Pentium 133/166 boxes destined for retirement.

Software development

Over 400 VM guests and 16 LPARs on the S/390 are being used for product development of DB2 and WebSphere Commerce Suite products.

IBM Internal Open Source Bazaar

The IBM Internal Open Source Bazaar is an internally deployed open source project development portal based on SourceForge to promote open source style development inside IBM. It is an opportunity for IBM developers to “share” their code with other IBM developers around the world. It is hosted on a Linux-based server. It is a service that offers easy access to the best in CVS, mailing lists, bug tracking, message boards/forums, task management, site hosting, permanent file archival, full backups, and total Web-based administration.

IBM Research facilities

The following projects are examples of where IBM Research facilities are using Linux-based servers.

Web Fountain Data Mining

This project has been developed into an eUtility service that is hosted on a multi-user infrastructure platform hosted at the Almaden Research Center, San Jose, California. This service uses an advanced data mining application that provides a revolutionary qualitative and quantitative analysis of the World Wide Web, intranet data, and other sources such as third party data sources. There is a development environment of 280 and a production environment of 358 Intel based servers running Linux, and growing rapidly.

Blue Gene® super-computer

This an advanced research project that is being developed by IBM and many research centers around the world to provide TeraFLOPs computing power running Linux on a three-dimensional mesh of interconnected processors. At the time this redbook is being written, around 10024 Linux servers are being used as a computing node in one machine, managed by a pSeries 690 running Linux.

This is an advanced super computing project that focuses on the modeling of the life science process of protein folding.

Oceano Web hosting utility

The IBM Research Division has developed hosting utility solution that can dynamically adjust capacity in response to workload across multiple commercial customers. This involves using a clustered set of Linux servers that can be dynamically reconfigured between customers. This has become the basis of IBM's on demand service offerings.

intraGrid

IBM grid computing plays a major role in e-business on demand™ and is an important cross-IBM play with benefits for all business units. Grid computing can bring tremendous productivity and efficiency to customers facing the challenges of an on demand world. The IBM community can find useful information about the nature of grid computing and the opportunities that grid can bring to customers and IBM alike.

IBM has available for customers and partners a full demo platform running mainly Linux and other supported operating system to help them to adopt this new concept of distributed computing.

intraGrid includes two objectives:

- ▶ The creation of an experimental grid computing infrastructure within IBM based on the Globus toolkit
- ▶ The creation of an IBM community interested in and participating in grid computing activities

The intraGrid Web site is aimed at allowing the sharing of resources, ideas, problems, and solutions. It is also aims to foster contributions to the community at large of grid demos, grid tools, grid applications, and grid information.

Within the scope of the intraGrid initiative, the Globus Toolkit 2.0 was made available as RPMs for various IBM platforms in 2002.

In February 2003, IBM began deploying a new set RPMs based on the Globus Toolkit 2.2.4 with additional customization and additions specific to the intraGrid. Currently, the RPM for RH Linux 7.1 or higher for x86 is available. Shortly, additional RPMs will be made available for these platforms:

- ▶ Linux on ppc (pSeries and iSeries)
- ▶ Linux on zSeries
- ▶ Linux on IBM @server running Opteron processors
- ▶ AIX on pSeries and on OS/400® Portable Application Solutions Environment (PASE) (iSeries)

This infrastructure is provided mainly for experimentation with and prototyping of grid technologies. We encourage the deployment of pilot applications on the IBM internal grid infrastructure. Such deployment and use are essential to test, evaluate, and tune grid technologies and tools.

Note: You can learn more about IBM grid computing solutions at the grid computing portal Web site at:

<http://www.ibm.com/grid>

Client for e-business

IBM offers their employees to use Linux on their desktop. For this, there is a Linux image, based on Red Hat 9, that has been modified to meet IBM security requirements. It also offers better support of the IBM line of desktops and notebooks.

This image provides all the productivity tools that any IBM employee needs to work, including:

- ▶ Lotus Notes 5 client for Linux
- ▶ SameTime client for Linux
- ▶ Microsoft compatible office suite
- ▶ Microsoft documents viewer

1.3 Linux: The operating system

Over the past few years, the Linux operating system has become a real and viable alternative for PC users as well as corporate servers and users. Linux delivers the power and flexibility of a UNIX server or desktop. It also provides a set of utilities, Internet applications, and a fully functional desktop interface.

The Linux operating system has become a server platform for powerful Internet and many other applications. Linux is capable of running from corporate Web, File Transfer Protocol (FTP), file and printer servers to wide-area information server (WAIS) Web sites or even corporate database servers, with real-time clusters or high performance computing clusters.

Linux is a fully functional operating system similar to a UNIX system. It has all the standard features of enterprise UNIX systems. Linux provides an easy way to migrate from other UNIX flavours offering the same capabilities and a similar way to work. Management of the command structure is enabled through *shells*. Enhanced graphical environments and desktops are also available.

Linux provides four fundamental components:

- ▶ **Kernel:** This provides low-level system control and interfaces, program and hardware device management, and an abstraction layer for the user level.
- ▶ **init:** This is the main program executed after the kernel is loaded into memory. From init, you can execute different run levels on the user space, with different characteristics.
- ▶ **File structure:** This structure is organized on file systems that provide the interfaces and abstraction needed to work with data and files. Files are organized into directories with the disk hardware. Each directory may contain any number of subdirectories each holding files. This is a similar structure to classical PC and UNIX operating system file structures. Linux supports the majority of the existing file systems, natively or through additional kernel modules.
- ▶ **Applications:** Linux applications cover many categories. The list is ever-expanding as more companies embrace this technology. Through writing applications for Linux, development costs usually decrease, spawning new projects and increasing the number of titles overall. Across the many horizontal and vertical markets with a Linux presence, there are 11 categories: office productivity, Internet related, network and systems management, software development tools, security, database related, graphic manipulation, audio/video production, data management, accounting and finance, and publishing.

Table 1-1 illustrates the simplified Linux structure.

Table 1-1 Simplified Linux structure

Shell	Web server	Web browser	Desktop environment	Office suite	E-mail client	SQL database	Productivity applications
init							
Kernel							
System hardware							

In some solutions, typically with clusters or non-PC platforms, the Linux server may not need the traditional PC hardware BIOS. The hardware can be directly controlled by Linux. This provides phenomenal bootup times (three seconds is the current record) that can be exploited by embedded devices and real-time needs.

Linux has the same multi-user and multitasking capabilities as large UNIX operating systems. It provides the same level of system administration that you find on standard UNIX systems. Users can run several programs concurrently, and there is a separate level for user space than kernel space. You can create user accounts for different users, and define their access rights to the files and system components. Installation of new devices and network connection and control is also provided as standard in the Linux operating system.

As a development environment, Linux can provide a powerful set of development tools to create industrial-strength applications. The development toolset includes, by default, the *GNU C Compiler*. A full set of commercial development software is also available.

The Linux structure provides the ability for programmers to access the hardware of the computer and the networks to which it is connected. This is achieved by the provision of a *hardware abstraction layer* where programs can take advantage of hardware features through a standard applications programming interface (API).

Linux programs can be portable to and from other versions of UNIX systems. Linux can use ANSI C language, combined with one of several portable graphical user interface (GUI) toolkits. Many other programming languages are available. These programs can be written for both UNIX systems and Windows or other servers if they use standard calls for standard libraries.

1.3.1 Linux BIOS management

In a cluster environment, the ability to quickly reconfigure hardware is particularly desirable. Linux can help with this task because it has the ability to boot without a BIOS. Current PCs used as cluster nodes depend on a vendor-supplied BIOS for booting. The BIOS, in turn, relies on inherently unreliable devices such as floppy disks and hard drives to boot the operating system. In addition, current BIOS software is unable to accommodate non-standard hardware, making it difficult to support experimental work. The BIOS is slow, often erroneous, redundant, and most importantly difficult to maintain. Imagine walking around with a keyboard and monitor to each of the 128 nodes in a cluster to change one BIOS setting.

LinuxBIOS loads the Linux kernel from NVRAM and essentially requires no moving parts other than the fan. It performs a minimal amount of hardware initialization before jumping to the kernel start and lets Linux do the rest. As a result, it is much faster (current record three seconds), which has sparked interest in the consumer electronics community as well. Moreover, updates can be performed over the network.

Using a real operating system to boot another operating system provides much greater flexibility than using a simple netboot program or the BIOS. Because Linux is the boot mechanism, it can boot over standard Ethernet or over other interconnects such as Myrinet, Quadrics, or SCI. It can use SSH connections to load the kernel, or it can use the InterMezzo caching file system or traditional network file system (NFS). Cluster nodes can be as simple as necessary, perhaps as simple as a central processing unit (CPU) and memory, no disk, no floppy, and no file system. The nodes will be much less autonomous, making them easier to maintain.

The LinuxBIOS project: LinuxBIOS is an open source project. It aims to replace normal BIOS with hardware initialization and a compressed Linux kernel that can be booted from a cold start. The project began as part of clustering research work in the Cluster Research Lab at the Advanced Computing Laboratory at Los Alamos National Laboratory. The motivation for the project was for the operating system to gain control of a cluster node at power on. Other benefits of LinuxBIOS include needing only two working motors to boot (CPU fan and power supply), fast boot times (current fastest is three seconds), and freedom from proprietary (buggy) BIOS code. Overall, the benefits are numerous, and helped to gain support from many vendors in the high-performance computing and embedded computing markets.

For more information about LinuxBIOS, refer to the following Web site:

<http://www.linuxbios.org/index.html>

1.3.2 Where Linux fits in

The adoption of Linux has been a rapid and changing process. It can be difficult to determine where Linux fits into the modern IT world. IBM is not in the early adopters or first deployments phase, but starting to implement Linux at the enterprise-level, enabling bigger and better capabilities. Figure 1-2 shows the modern IT spectrum, from the operating systems to the client interface. Linux plays a role in almost every area, and is quickly making inroads where traditional operating systems ruled for years.




Figure 1-2 The IT spectrum: Logically and physically

1.3.3 Working with other operating systems

Linux is an open operating system that tolerates, in a powerful way, working together with or inside of another operating system on a system. It is frequently used to populate virtual partitions on larger, non-Intel based systems such as the iSeries, pSeries, and zSeries servers. The virtual or logical system partition provides three key features:

- ▶ **Server consolidation:** An organization with existing Linux or UNIX servers can save hardware and personnel costs by consolidating those servers into virtual partitions on an existing IBM @server platform.
- ▶ **High availability:** Multiple virtual partitions can provide a level of availability and application-based load balancing while retaining the management benefits of consolidated hardware.
- ▶ **Infrastructure:** Linux virtual partitions can provide security and infrastructure services to other partitions that share the same virtual network.

1.4 Infrastructure

Linux enters the mainstream markets by providing critical infrastructure services.

Web serving

The combination of Linux and the Apache Web server, or other Linux supported Web servers such as Zeus or Netscape, offers an attractive package for customers. It provides a low-cost, flexible solution originally for static Web sites, with over 30% of the world's Web sites running this combination. The demand is now moving toward a more dynamic approach with Web sites that users can interact with, and that support high transaction rates.

File and print serving

One of the basics for Linux implementation is the provision of inexpensive facilities such as file and print services. Linux offers a rapid return on investment (ROI) in this part of the infrastructure space. The management capabilities and low cost make this an easy solution to justify. Also, this is an important environment, but it does not typically have the operational importance of line-of-business applications. It is a relatively safe place for businesses to test this new technology.

Domain name server (DNS) and DHCP

As a UNIX-like operating system, Linux is well proven at hosting Berkeley Internet Name Daemon (BIND) name servers and Dynamic Host Configuration Protocol (DHCP) services.

Router

Linux is capable of advanced routing using inexpensive commodity hardware. Also, some router vendors have chosen Linux to be their embedded operating system.

Firewall and Intrusion Detection Services (IDS)

Linux has been a popular provider of firewall and IDS services. Because of the advanced configuration and customization options, along with a small memory footprint, Linux has been an ideal solution for many organizations who want to avoid proprietary solutions.

1.5 Data storage

The explosive growth in data storage needs has driven Linux into a premier position as a low-cost, efficient, and highly customizable platform to handle

storage management. In particular, Linux is being used in Internet Small Computer Systems Interface (iSCSI) environments. iSCSI works by attaching SCSI-based storage resources directly to an Ethernet network and is a low-cost alternative to Fibre Channel solutions. In this environment, Linux operates as a storage node, managing RAID hardware, and providing a high level of network performance.

Refer to Chapter 8, “TotalStorage and Linux” on page 237, for more information about IBM TotalStorage solutions.

1.6 Middleware

Linux is currently seeing large growth in the middleware arena. Middleware is called the “glue” that holds IT applications together. It is typically divided into the following four areas.

Database

Linux is quickly growing as a preferred platform for databases because of its low cost, stability, and extensibility. It also offers enterprise-level features such as load balancing and high availability clusters. The two most common database implementations running Linux are DB2 and Oracle 9i RAC.

Note: IBM DB2 Universal Database™ (UDB) provides enterprise-level database features on Linux. For more information, see:

<http://www.ibm.com/software/data/db2/linux/validate/>

Oracle 9i Real Application Cluster (RAC) also provides enterprise-level clustered database features. For more information, see:

<http://www.oracle.com>

Cross-system support

Most products with the “middleware” label are really software platforms that provide integration, transaction, development support to tie various IT services together. Linux is the operating system of support for middleware because its open and flexible nature works well in the heterogeneous environment of modern IT networks. IBM offers two leading middleware examples:

- ▶ WebSphere
- ▶ MQSeries®

Note: The IBM WebSphere platform provides integration, transaction, and development support services while taking advantage of the open nature of Linux. To learn more about WebSphere, see this site:

<http://www.ibm.com/software/info1/websphere/index.jsp>

Knowledge discovery and collaboration

Increasingly, a major task of middleware is to organize the information flowing through a network, and not to just manipulate and store it. Knowledge discovery (KD) refers to the integration of enterprise content management (ECM) and knowledge management (KM). It entails merging multiple technologies that originally (and individually) were responsible for capturing, locating, organizing, and delivering information.

Collaboration includes all technologies that provide integrated messaging, calendaring, and task management. The open heritage of Linux provides a flexible platform in which to integrate the various knowledge discovery and collaboration tools.

Note: IBM provides knowledge collaboration products through its Lotus line. Follow this link for more information about Linux and Lotus:

<http://www.lotus.com/>

Systems management

Systems management has increasingly become centralized and heterogeneous. With an decreased margin for downtime, middleware must incorporate the ability to monitor all systems at all levels of abstraction. Because Linux is open in nature, it allows a higher level of access to its internal workings than closed or proprietary systems offer.

A management infrastructure built on Linux also takes advantage of its UNIX system heritage, and many of the third-party and open source products that can provide additional functionality.

Note: IBM provides systems management middleware through its IBM Tivoli software family. You can learn more about IBM Tivoli on the Web at:

<http://www.ibm.com/tivoli>

1.7 Application solutions

Application solutions refer to those applications that are critical for day-to-day operation of an organization. They usually encompass Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), Supply Chain Management (SCM), and Business Intelligence (BI) products. Some Linux-capable solution providers are introduced in the following sections. They include SAP, J.D. Edwards, and eOne Group.

Overall, there is a wide range of support for Linux across the entire spectrum of IT system requirements. Table 1-2 summarizes some of the key elements and the applications that fill them.

Table 1-2 Business areas and solutions

Infrastructure solutions	Vendor or solution
Database and Application Development Tools	DB2 UDB, Everyplace, Oracle, Informix (Retail/Distribution) Sybase (FSS), Rational
Application Server	WebSphere Application Server, Apache, Domino
Network and System Management	Tivoli, VERITAS, BMC, Trustix, Compuware
Security and Firewall	Checkpoint, Symantec, CISCO, Trustix, SuSE Firewall, Tivoli
Clustering (High Availability)	Steeleye, VERITAS, PolyServe, Legato
Server Consolidation	VMware, Samba
Mail and Messaging	Sendmail, Bynari, Domino, Caldera VMS, SuSE eMail Server
e-business solutions	
E-commerce	WebSphere Commerce Suite, eOneGroup
Accounting and CRM	ACCPAC, Appgen, Sage
Business Intelligence and Data Warehouse	SAP BI, Brio, Crystal Decisions, Dimensional Insight, Hyperion
ERM/CRM/SCM (cross industry)	mySAP.com, QAD, BaaN Invensys
Retail Distribution	Relavis, MarCole, 360 Commerce, Caldera
Finance and Banking	Axiom, Sanchez, Rueters, Intelligent Markets, JD Edwards CRM for Financial Services, SUNGARD
Industrial (EDA)	Synopsis, Cadence, LSDC, Mentor Graphics

SAP

SAP supports Linux versions of their ERP products, but restricts support to specific hardware platforms. This Web site provides a list of SAP-certified Linux hardware platforms:

<http://sap.com/solutions/technology/linux/platforms/index.asp>

J.D. Edwards

J.D. Edwards has announced that they will offer their Financial Services CRM tool with the IBM WebSphere Application Server running on a Linux and xSeries server.

The following Web site provides announcement details:

<http://www.jdedwards.com/pressreleases/>

eOne Group

eOneCommerce is 100% platform independent. You can integrate it with IBM hardware, software, database, and operating system. It can run in a Linux LPAR or under a native operating system.

1.8 Clients

While desktop versions of Linux traditionally trailed Microsoft products in ease-of-use, that is rapidly starting to change. There is an increasing use of Linux as a client interface in all areas of computing. The following sections provide specific examples.

Thick clients

Thick clients include desktops and laptops that run Linux as their primary operating system. However, it is still common (especially in development groups) to see dual-booting machines to allow some Windows compatibility.

Thick clients are prevalent in the business and technical support organizations.

Thin clients

High maintenance and labor costs have contributed to the trend of replacing thick clients with thin clients.

Linux makes an ideal thin client platform due to the limited memory requirements and the broad embedded processor support. The highly customizable nature of

Linux allied with centralized control makes the thin client an attractive business proposition.

Note: The IBM NetVista™ line offers several Linux-based thin client solutions with detailed information. You can learn more on the following Web site:

<http://www-3.ibm.com/pc/support/site.wss/document.do?lnidocid=MIGR-48NT8D&selectarea=SUPPORT&tempselected=5>

Personal digital assistants

Linux has been a popular choice for embedded devices for all types. However, commercial Linux-based personal digital assistants (PDAs) only became available in the last few years.

Application portability is cited as the major benefit, because users can run the exact same applications on their PDAs as they do on their desktops. An excellent example of a Linux-based PDA is the *Paron MPC*. The Paron was developed primarily for vertical markets that need secure, encrypted access to a Bluetooth-enabled infrastructure like banking, inventory, and government. Paron MPC offers a finger-print recognition chip as part of its security features.

IBM co-developed the Paron MPC with CDL. Refer to the following sites for details:

<http://www.linuxdevices.com/articles/AT7145548309.html>
http://www.cdlusa.com/products/paron_mpc.shtml

Other companies, such as Sharp and G-Mate, also develop Linux-based PDAs such as Sharp's Zaurus family and G-Mate's Yopy. You can learn more about those products online at:

<http://www.zaurus.com>
<http://www.yopy.com>

1.9 The Linux distributions

What is usually referred to as *Linux* is more accurately called a “distribution” or “distro” for short. Technically, Linux is only the core or kernel of the operating system. Linux distributions are prepackaged collections of software, generally comprising the kernel and various other packages.

Important: The following descriptions of Linux distributions are a point in time and an overview. The reader should always visit the distributor's website for the most up to date information about the product.

The major distributors include a graphical desktop environment and an installation program to help you get started. The kernel (for which the source code is freely available) is managed by Linux creator Linus Torvalds, and a leadership team of cooperative and distributed work. Many developers contribute to updating the kernel with new features as well as device drivers. The packages that can be added to the operating system include compilers, shells, and other applications. Not all of these packages may be available as open source. You can learn about open source in 1.9.3, "What is open source" on page 29.

A Linux distribution includes the Linux kernel plus utilities, programming tools, window managers, and other software that make up a full operating system. Distribution companies such as Red Hat, SuSE, and Turbolinux, and non-commercial organizations such as Debian build their own distribution around the basic kernel, which is the core of the distribution. Then they add utilities such as installation programs and package management tools. Next they package them on an image that can be presented on a CD-ROM/DVD or downloadable on the Internet. The kernel code in each distribution may be different since distributions pull Linux versions. The Linux Standard Base project attempts to standardize the aspects of a distribution to minimize the changes that may affect portability.

A *distributor* (such as Red Hat or SuSE) is a company who packages a Linux distribution for commercial purposes. They technically do not resell the software license, but rather sell a subscription service for support and maintenance. Generally, these are one-year subscriptions. Distributors also sell Linux consulting and integration services. They may sell applications or middleware in addition to the Linux operating system.

A search using Google (December 2003) produces a hit about the occurrences of the distributions on the Web. For each distribution, Table 1-3 shows the number of occurrences of its name on the Web (more user-friendly), in the Newsgroups (more technical), and the ratio of both, which we refer to as the "talk factor". An important talk factor means that the distribution is very active. It also means that more people need to talk about it to make things work as they want often because they want more than the standard user.

Table 1-3 Presence on the Internet of various distributions

Distribution	Web references	Newsgroup references	Talk factor
Red Hat	3660000	85500	0.023
SUSE	2950000	64900	0.022
Debian	2880000	227000	0.079
Mandrake	1940000	46100	0.024
Slackware	887000	36600	0.041
Gentoo	836000	17400	0.021
Knoppix	398000	13100	0.033
Conectiva	306000	2260	0.007

1.9.1 Inside the distributions

This section introduces you to the Linux distributions. It answers the question: “What does a Linux distribution offer?”

How do you know which is the right distribution for your environment? Hundreds of Linux distributions are currently available. The move into enterprise class software drives and separates the distributions. Examples of enterprise class distributions include SuSE Linux Enterprise Server (SLES) and Red Hat Linux (RHEL).

Not all distributions are supported on all IBM @server platforms. In general, IBM supports those distributions that provide the enterprise class of service. These distributions cost more, but include more enterprise class features and multiple levels of support from bug fixes to 24-x-7 support.

Enterprise class distributions include the following features:

► **New offerings:**

- Optimized release cycles of 12 to 18 months, three to five or more year version support
- Independent software vendor (ISV) certification platform
- Increased commitment to maintenance and support
- Enterprise rich functionality such as multi-threaded input/output (I/O), improved SMP usage, and more

► **Customer value:**

- *Reliability*: Longer integration and testing cycles
- *Scalability*: Improved SMP; functions such as Piranha from Red Hat for load balancing, and SuSE backport of 2.6 scheduler for scalability and performance

- *Manageability and service*: Improved RAS function, 24-x-7 global enterprise class support offerings

► **Implementation:**

- Red Hat Linux
- SuSE Linux Enterprise Server 8 powered by UnitedLinux
- Turbolinux Enterprise Server 8 powered by UnitedLinux

This section presents some of the most popular distributions, in particular those that run on IBM *@server* independent software vendor platforms and give you a preview of each. Then it highlights other players in the marketplace. That includes distributions that are not as well known, but have an impact in specific geographies or are suited to a different audience.

Application compatibility

In the vast majority of cases, Linux applications are compatible with all distributions of Linux, which accounts for the aphorism “Linux is Linux is Linux.” When a new Linux kernel is released, it is uploaded to the main Linux kernel site:

<http://www.kernel.org>

The distribution companies then package the kernel and provide it as an update on their own Web sites. Distribution vendors take the kernel as is. This includes all changes and fixes that are contributed by members of the development community. They may also add updates or patches to add specific features to their business needs. The speed of distribution release has slowed considerably in the last few years. At one point, new distributions were released about twice a year. With the emphasis on the enterprise class distribution and wide range of products running on or under Linux, the pace of releases has slowed considerably.

Each distribution categorizes its products uniquely. For example, Red Hat divides its products into enterprise and community versions. SuSE has two divisions: business and home users. Turbolinux showcases its products in enterprise and desktop. Where possible in this chapter, we separate the products into business and workstation markets.

As new hardware is announced and IBM Linux distribution partners release new versions, it is a good idea to check the status of IBM support. The following Web links can help you find the latest status for each IBM *@server*. The IBM development team also writes installation instructions for each system, as it leaves testing, to document any recommendations for running versions of Linux operating systems.

For the Linux version that the IBM @server platforms support, consult these sites:

► xSeries

<http://www.pc.ibm.com/us/compat/nos/matrix.shtml>
<http://www-1.ibm.com/servers/eserver/linux/xseries/>

► pSeries

<http://www-1.ibm.com/servers/eserver/pseries/linux/index.html>

► iSeries

<http://www-1.ibm.com/servers/eserver/iseries/linux/>

► zSeries

<http://www-1.ibm.com/servers/eserver/zseries/os/linux/index.html>

You can find the IBM supported options for xSeries servers, sorted by system and operating system, on the following site:

<http://exist.raleigh.ibm.com/nosmgr/optlist/sprec.asp>

IBM collaborates with Red Hat, SuSE, and Turbolinux, certifying hardware through the Linux distribution partner certification programs. The results are posted within the Linux distribution partner Web sites:

► Red Hat:

<http://hardware.redhat.com/hcl/?pagename=hcl>

► SuSE:

http://hardwaredb.suse.de/index.php?LANG=en_UKd

► Turbolinux:

<http://www.turbolinux.com/hcl/TTlist.html>

1.9.2 What is common among the distributions

There is much in common among each distribution. For example, each distribution comes with two major desktop GUI interfaces: the K Desktop Environment (KDE) and GNOME. To the end user, they are functionally equivalent. Choosing one over the other is a matter of personal preference. Both of these Windows desktop GUIs function well as graphical interfaces. For developers, the difference is the set of tools and libraries that each option offers. KDE uses the Qt C++ cross-platform GUI toolkit for KDE development, where GNOME uses the Graphics Tool Kit (GTK) and other tools and utilities.

Each distribution provides an installation program. Most of them are GUI programs that also support execution in text-mode. An office suite is usually included with each distribution such as Star Office, OpenOffice.org, or a distribution-created office suite. Most distributors also contain a variety of utilities that include configuration and installation utilities, firewalls, Web browsers, and Web server programs.

The distributions include differences in the versions of the kernel, compilers, libraries, languages packs, and adherence to standards. The UnitedLinux initiative seeks to answer some of these issues with standards. Red Hat is also promoting such a model with their Open Source Architecture project. Support can also vary from distributor to distributor, for example, response times to problems, sales and support coverage by geography, etc.

1.9.3 What is open source

Open source is often thought of as software that is given away, or free to be redistributed. In reality, open source is software whose *source code* is available free of charge to the general public. The rationale for this concept is that a larger group of programmers, who are not concerned with proprietary ownership or financial gain, will produce a more useful and bug-free product for everyone to use.

Here are three examples that illustrate how the Internet contains open-source software in heavy commercial use:

- ▶ **Apache:** Runs over 70% of the world's Web servers
- ▶ **Perl:** The engine behind most of the *live content* on the World Wide Web
- ▶ **BIND:** The software that provides the domain name service (DNS) for the entire Internet

Note: Ten items comprise the Open Source Definition. For a detailed discussion, see the Open Source Initiative (OSI) Web site at:

<http://www.opensource.org>

You can read more about open source in 2.2.2, “The compromise of open source software” on page 50.

Linux is an example of open source software. Linux programmers, working over the Internet, are the community that improves, adapts, and fixes Linux with amazing speed. This rapid evolutionary process frequently produces better software than the traditional development model.

Linux: An example of open source use

Linux is written and distributed under the GNU GPL. The GNU GPL is an example of certified open source software. As is the case for open source, Linux's source code is freely distributed and available to the general public.

The GPL ensures the following four points:

- ▶ Freedom to distribute copies of free software (and charge for this service if you want)
- ▶ Receive source code or obtain it upon request
- ▶ Software modification or use in new, free programs
- ▶ Information about the user's right to perform these actions

The entire GNU GPL script appears at this Web site:

<http://www.gnu.org/copyleft/gpl.html>

1.9.4 The Linux Standard Base project

The Linux Standard Base project is an earlier attempt designed to bring the array of Linux distributions into a common core. It is a workgroup of the Free Standards Group.

Supported by the development community and IT industry leaders, the Free Standards Group is an independent, non-profit organization dedicated to accelerating the use and acceptance of open source technologies through the development, application, and promotion of standards. These standards include common behavioral specifications, tools, and APIs to make development easier across Linux distributions.

The majority of Linux distributions that are available on the market tend to meet the Linux Standard Base definitions and requirements to make a unique common base for all distributions.

For more information, see the Linux Standard Base home page:

<http://www.linuxbase.org>

1.10 Red Hat

Founded in 1994, Red Hat provides the most popular distribution in North America. It is well known for its Red Hat Package Manager (RPM) system for

installing and maintaining its components. It was one of the first package management tools available for Linux.

Red Hat offers development, deployment, and management tools for Linux and open source environments, ranging from embedded devices to enterprise servers. The Red Hat Linux distribution is also the basis for several other Linux distributions.

For an informed overview of the Red Hat products read the following white paper available at:

<http://www.redhat.com/whitepapers/rhel/RHEL3FamOverWPS0.pdf>

Red Hat Linux operating system products

Red Hat changed their release model, and now they work in two different offerings for different users. The new releases include:

- ▶ Red Hat Enterprise Linux , which offers three different products that are oriented to the enterprise market:
 - Red Hat Enterprise Linux AS
 - Red Hat Enterprise Linux ES
 - Red Hat Enterprise Linux WS
- ▶ Red Hat Fedora project: This distribution is completely free to download, open source, and community developed.

The releases that were available at this time this redbook was written are:

- ▶ Red Hat Enterprise Linux 3
- ▶ Red Hat Linux 9

Note: Business, government, and individuals looking for a stable, supported Linux should research *Red Hat Enterprise Linux*. Developers and technology enthusiasts should investigate the *Fedora Project*, a proving ground for new Linux and open source technology.

For a complete comparison between the three Red Hat product lines, see this document:

<http://www.redhat.com/software/rhel/comparison/>

For additional IBM information, news, products, press releases, services, education, and support on Red Hat, see:

http://www-1.ibm.com/linux/va_4010.shtml

For additional information about Red Hat, visit their Web site at:

<http://www.redhat.com>

1.11 UnitedLinux

On 30 May 2002, Conectiva, SCO Group, SuSE, and Turbolinux announced that they would join forces to create a single binary-compatible Linux distribution, sharing development costs and future royalties. UnitedLinux was formed as a consortium, and the UnitedLinux product UL 1.0 was ultimately announced in November of 2002. By 4 December 2002, SCO unveiled SCO Linux 4.0, and SuSE unveiled its SuSE Linux Enterprise Server 8. These were the first Linux distributions on the market based on the UnitedLinux 1.0 code. Turbolinux released its Enterprise Server 8 version in Japan.

UnitedLinux is a standards-based Linux operating system. It targets the business user. It seeks to streamline Linux development and certification around a global, uniform distribution of Linux. It is certified to work across the both AMD and Intel 32-bit x86 lines; the 64-bit x86 Athlon, Opteron, and Itanium families; and the complete IBM *@server* brand.

While each distributor sells the same core UnitedLinux distribution, they can add additional programs and utilities, provided that application portability is preserved. Each distributor sell its products under its own brand name “Powered by UnitedLinux”.

For more information, see the UnitedLinux Application Brief on the Web at:

<http://www.ibm.com/linux/UnitedLinux-Application-Brief.pdf>

You can also read more about UnitedLinux in their technical white paper at:

<http://www.unitedlinux.com/pdfs/whitepaper4.pdf>

Major companies and ISVs can now work with two main Linux versions (the other being Red Hat) instead of five or more major versions. This simplification saves companies from porting applications across multiple distributions. In its first month, over 15,000 downloads of the UnitedLinux beta were initiated.

1.11.1 What’s in the distribution

UnitedLinux contains a common kernel, APIs, file systems, and installation routines for business Linux. Developers and customers obtain UnitedLinux through one of UnitedLinux’s supporting vendors. That is, you order it from either Conectiva, SuSE, or Turbolinux. Regardless of the vendor name on the box, you receive UnitedLinux. This comes in the form of an installation CD or CDs. Each

company then adds their own software packages, which come on separate CDs. Figure 1-3 shows the architecture of UnitedLinux.

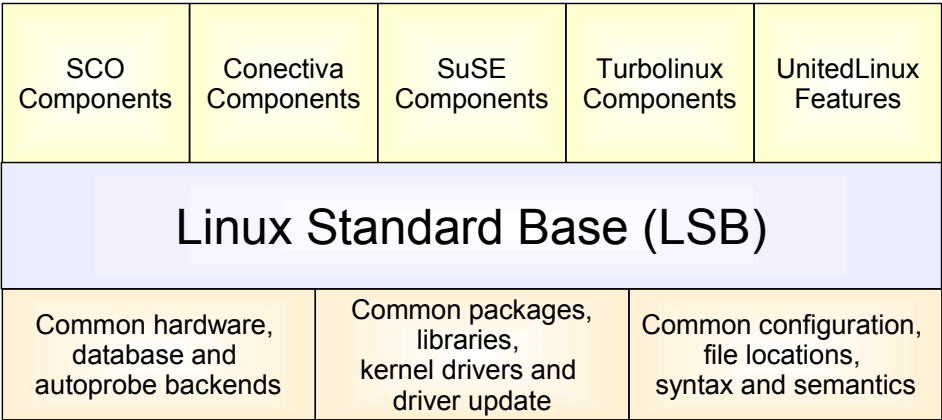


Figure 1-3 UnitedLinux architecture

UnitedLinux is a General Public License Linux. The source code is free, but the binaries (executable code) are not. For example, using the UnitedLinux source code, developers can create their own binaries (compiled source code ready for execution), but they will not be allowed to use UnitedLinux branding or UnitedLinux support.

In addition, UnitedLinux supports the Linux Standards Base (LSB) described in 1.9.4, “The Linux Standard Base project” on page 30. You can view UnitedLinux as an implementation of the LSB and the Li18nux and GB18030 standards. UnitedLinux also has international support for the English, German, French, Italian, Japanese, Korean, Portuguese, Spanish, Simplified Chinese, and Traditional Chinese languages. For more information, see:

<http://www.unitedlinux.com>

The standards UnitedLinux requires of its members

UnitedLinux adheres to Linux and industry standards, including:

- ▶ **File System Hierarchy Standard (FHS):** A set of requirements guidelines for file and directory placement under UNIX-like operating systems
- ▶ **Linux Standard Base:** Standards that increase compatibility
- ▶ **LI18NUNIX:** Internationalization of a core set of APIs and components

- ▶ **Scalability:** To take advantage of larger and more complex systems and handle new categories of applications:
 - A complete set of software and tools to build server farms for workloads that would otherwise be unmanageable for a single machine
 - Scheduler enhancements to improve process scheduling on SMPs and to prevent the scheduler from becoming a bottleneck
 - Asynchronous input/output to minimize waiting on I/O in large, busy systems
- ▶ **High availability:** The ability to avoid downtime, including:
 - POSIX-compliant event logging and notification capability
 - Dynamic probes: Dynamic insertion of breakpoints in code for debugging
 - Non-disruptive and tailored dumping of system data
 - Toolkit to record and trace system events
 - Hotplug PCI support
- ▶ **Security:** Kerberos support, basic firewall support, and a consolidated set of community security enhancements known as *Bastille*
- ▶ **File systems:** Support for the journaling file system (JFS), and the popular community file systems Reiser File System (ReiserFS), XFS, and the ext3 file system
- ▶ **Network/storage/device management:** Supports logical volume manager (LVM) and Enterprise Volume Management System (EVMS) to manage storage
- ▶ **Platform support and interoperability:** Hardware capability across all relevant platforms and architectures, including Intel (32 and 64-bit), AMD, PowerPC® (iSeries and pSeries), and the zSeries mainframe
- ▶ **Development environment:** A development environment for ISVs that includes all the compilers for libraries, sources, text editors, graphical user interface support, and other tools to enable the building of applications for UnitedLinux

For more information, see the UnitedLinux white paper at:

<http://www.unitedlinux.com>

1.12 Conectiva

Conectiva, a UnitedLinux partner, is a Brazilian Linux distribution with several titles in Spanish and Portuguese. This Linux distribution is created specifically for Latin America. Conectiva Linux is aimed at both desktops and servers. It is

RPM-based and APT-enabled. It can be installed using several preconfigured profiles, or it can be completely customized for specific tasks or power users.

Conectiva includes over 1,000 applications including Star Office, Netscape (in native languages), games, image manipulation applications, and other utilities. It also includes a wide array of commercial software, such as Oracle, VariCAD, and ViaVoice®.

The Server Edition comes with software for network administration, e-commerce, and support for RAID and clusters. The current version includes native support for the standard network protocols including IPv6, NIS, and LDAP. Additional features include remote boot, Webmail, load balancing, RAID, FAX, Groupware, virtual file system (VFS), security auditing, and mail and Web servers.

Conectiva contains two different categories of profiles, one for servers and one for work stations. The file systems that are supported include Ext2fs, Reiserfs, NTFS, FAT, and FAT32.

Conectiva ships with Distributed Replicated Block Device (DRBD), a high availability kernel module that mirrors hard drive contents over a network. A high availability utility, the *heartbeat package*, monitors machines within a cluster, and notifies the setup when one of the machines within the cluster dies, allowing the box with the mirrored drive to take over.

You can learn more at Conectiva's Web site:

<http://www.conectiva.com>

1.13 SuSE

Founded in 1992, SuSE Linux is based in Nuremberg, Germany, and historically has maintained a high marketshare in Europe. The latest versions have a new graphical installation tool to simplify the installation. The advanced features include hardware acceleration for 3-D graphic chips and a journaling file system. SuSE also develops and supports a series of X Servers for newer, high-end graphic cards. The installation is available in German, English, Italian, and French. SuSE Linux comes with a large group of utilities and applications.

SuSE is available in business versions. These include SuSE Linux Enterprise Server (SLES); servers for zSeries, xSeries, pSeries, and iSeries; and versions that support Alpha, PowerPC, and Intel platforms.

Important: In December 2003, Novell announced the acquirement of SuSE. This is part of Novell's new strategy on Linux. Novell also recently acquired Ximian, which is a company focused on Linux desktop solutions. Due to these acquisitions, some product names in this edition of this Linux Handbook may change.

For more information on Novell and SuSE visit:

<http://www.novell.com/news/press/archive/2003/11/pr03069.html>

1.13.1 The distribution

SuSE divides its products into two categories: business customers and private customers. The following sections describe the major offerings and their features.

Solutions for business customers

SuSE Linux Enterprise Server is a server operating system for deployment in IT environments of all sizes. It is available for multiple hardware platforms:

- ▶ SuSE Linux Enterprise Server (SLES) 8
- ▶ SuSE Linux Standard 8

SuSE also supports IBM mainframe capabilities in S/390 and zSeries including channel architecture, virtualization, crypto hardware support, and hipersockets.

Other business solutions include:

- ▶ SuSE Linux Openexchange Server 4.1: A groupware and e-mail server
- ▶ SuSE Linux Desktop

Solutions for private customers

For private customers, SuSE offers two products:

- ▶ SuSE Linux 9.0 Professional
- ▶ SuSE Linux 9.0 Personal

For additional IBM information, news, products, press releases, services, education, and support on Red Hat, see:

http://www-1.ibm.com/linux/va_4011.shtml

For more information, visit the SuSE Web site at:

<http://www.suse.com>

1.14 Turbolinux

Turbolinux incorporates a suite of Linux products for both the workstation and server. It is particularly used in the Asia Pacific markets, having both Japanese and Chinese language integrated versions on top of the unified code base that is sold as the international (English) version.

Turbolinux offers Linux solutions geared toward corporate needs and home desktop needs. Turbolinux also offers a unique clustering solution that allows for the construction of highly available, and scalable networks based on low-cost commodity components.

1.14.1 What's in the distribution

Turbolinux provides ten business solutions:

- ▶ Turbolinux Enterprise Server 8 powered by UnitedLinux (IA32)
- ▶ Turbolinux Enterprise Server 8 powered by UnitedLinux (Itanium Processor Family) in process of release
- ▶ Turbolinux 8 Server (IA32)
- ▶ Turbolinux 8 for AMD X86-64
- ▶ Turbolinux 8 Workstation (IA32)
- ▶ Turbolinux Enterprise Server 8 for zSeries powered by UnitedLinux
- ▶ Turbolinux Enterprise Server 8 for iSeries powered by UnitedLinux
- ▶ Turbolinux Enterprise Server 8 for pSeries powered by UnitedLinux

Turbolinux 8 Server offers enterprise features. These include 64 GB memory support, large file support (LFS), which provides a maximum file size of 4 terabyte (TB) and 16 TB file system support, optimized scheduler, and LVM with disk volume management. It also includes journaling file systems, PXE support, Linux Kernel Crash Dump (LKCD) for problem analysis, and EVMS support.

For more information on Turbolinux Enterprise Server 8 read the datasheet at:

<http://www.turbolinux.com/products/tles8/tles8-datasheet-en.pdf>

Turbolinux Server for zSeries and S/390 is the first Turbolinux distribution optimized to run on IBM mainframe hardware platforms. It is a complete Linux server distribution. It includes the Linux operating system, database applications, an Apache Web Server, a Samba network file server, sendmail mail server, C programming language application development tools, and an FTP server, DNS, and Telnet.

Turbolinux Server for iSeries supports English, Japanese, Korean, and Simplified and Traditional Chinese with the unified code base. The Turbolinux

Server for iSeries distribution includes database applications, an Apache Web Server, a Samba network file server, sendmail mail server, C programming language application development tools, as well as an FTP server, DNS, and Telnet.

Turbolinux Server for pSeries supports English, Japanese, Korean, and Simplified and Traditional Chinese with the unified code base. Turbolinux Server for pSeries is a complete Linux server distribution that includes the Linux operating system, database applications, an Apache Web Server, a Samba network file server, and a sendmail mail server. It also includes C programming language application development tools, as well as an FTP server, DNS, and Telnet.

Turbolinux Server for Itanium extends the Turbolinux Server to Intel's 64-bit architecture. e-business, Internet, print, file, and database configuration options are available in the base package.

Turbolinux 8 Workstation includes e-mail, Web browsers, office and mail suites, PDA hot-sync capabilities, ADSL setup tool, MP3 and MIDI, and diverse hardware support including USB devices. In addition, Turbolinux 8 Workstation provides an ADSL interface for high-speed connections and extended support for various digital media including digital video data.

To learn more about Turbolinux, visit their Web site:

<http://www.turbolinux.com>

1.15 Other players

There are hundreds of Linux distributions. You can find between 150 to 300 documented summaries on some Web sites.

Several excellent resources list and summarize distributions. Among these, we recommend these sites:

- ▶ Linux Online, the main portal reference for Linux, and the beginner's starting point to more information about Linux
<http://www.linux.org>
- ▶ LinuxBasis, an interesting portal with information about distributions, how to's, and links to Linux information online
<http://www.linuxbasis.com/distributions.html>
- ▶ DistroWatch, a complete portal with information about all the Linux distributions available and updated daily

<http://www.distrowatch.com>

Due to the dynamic nature of Linux, we find continual changes as to which distributions are ranked in the top ten. If we do not mention your favorite distributor, it is not an oversight. Rather we highlight those that are currently most visible. The Linux world is in constant change, so we encourage you to investigate the resources yourself to find the latest information.

The other players we highlight include:

- ▶ Red Flag, the major distributor in People's Republic of China
- ▶ Yellow Dog, a Red Hat derivative that focuses on the PowerPC market, with a large South American install base
- ▶ Debian, a non-commercial distribution originally targeted for developers

Table 1-4 compares these players.

Table 1-4 Basic distributor information

Distribution name	Developed by	Platforms	Server, Workstation, Enterprise	Location on the Web
Conectiva	Conectiva Curitiba, Brazil	Intel	S - W	http://www.conectiva.com
Debian	Ian Murdock	Alpha, Intel, PPC, Sparc, Other	S - W	http://www.debian.org
Fedora	Linux Community; Red Hat sponsored	Intel	W	http://fedora.redhat.com http://www.fedora.us Note: Sites will officially merge at some point in 2004.
Red Flag	Red Flag Software Co., Ltd. Beijing, PRC	Alpha, Intel, PA, RISC, Sparc, S/390	S - W - E	http://www.redflag-linux.com/ eindex.html
Red Hat	Red Hat, Inc. Raleigh, NC	Alpha, Intel, Itanium/x86-64, AMD Opteron, zSeries, S/390, iSeries, pSeries, other	S - W - E	http://www.redhat.com
SuSE	SuSE Linux AG Nuremberg, Germany	Alpha, Intel, Itanium/x86-64, PPC, Sparc, zSeries, S/390 iSeries, pSeries, other	S - W - E	http://www.suse.com
Turbolinux	Turbolinux, Inc., Tokyo, Japan	Intel, Itanium, S/390	S - W - E	http://www.turbolinux.com
Yellow Dog	Terra Soft Solutions, Inc. Loveland, CO	PPC	S - W	http://www.yellowdoglinux.com

1.15.1 Yellow Dog Linux

Yellow Dog Linux 3.0.1 is a complete Red Hat and RPM-based distribution for the PowerPC architecture. Yellow Dog Linux version 3.0 is Terra Soft's most recent release, which offers a completely rebuilt Installer, 1300 packages on six CDs (three install, three source), a unified KDE and GNOME desktop environment featuring shared menus, applications, and user interface.

What's in the distribution

Yellow Dog Linux 3.0 includes the following components:

- ▶ 6 CDs (3 binary, 3 source)
- ▶ Anaconda installer engine
- ▶ Most advanced, stable kernels
- ▶ RPM Package Manager
- ▶ GNU LibC (glibc) and GNU Compiler Collection (gcc)
- ▶ XFree86, KDE, GNOME 2.2
- ▶ Support for NVidia and ATI Radeon graphics cards
- ▶ Mac-on-Linux enables both Classic and OS X to run on Linux
- ▶ OpenOffice.org
- ▶ Both “apt-get” and “yum” enable automated RPM updates.

Black Lab Cluster Management is Terrasoft's cluster product. Black Lab v2.2 is a cluster build and management application for High Performance Computing (HPC) clusters running Yellow Dog Linux. It offers single-click installation and configuration, automated updates through apt-get, and a graphical user interface.

You can learn more from the Yellow Dog Linux Web site:

<http://yellowdoglinux.com/products>

1.15.2 Red Flag Linux

Red Flag Software Co., Limited first appeared in August 1999, when it was created by the Institute of Software at the Chinese Academy of Sciences. Financial help came from government-owned Shanghai NewMargin Venture Capital. Red Flag came out with an integrated desktop office suite called *Chinese 2000*. It joined the GNOME foundation with the intent to localize GNOME for the Chinese population. It has training and certification programs such as Red Flag Certified Engineer (RCE) and several applications.

Red Flag Linux is the government-approved operating system for desktop PCs in the People's Republic of China. It won the Chinese government contract at the start of 2002 and the contract for the Beijing Municipal Government.

What's in the distribution

Red Flag provides many software products for enterprise computing, from embedded operating system made for embedded device, to cluster server products suitable for enterprise distributed computing environment, which creates the wide-range of solutions. Following the continuous supplement and development of product lines, Red Flag provides five main product lines in the enterprise computing market:

- ▶ Red Flag Linux Desktop series
- ▶ Red Flag Linux Server series
- ▶ Red Flag Application Server series
- ▶ Red Flag Solution series
- ▶ Red Flag Advanced Service series

Red Flag Server 4.0 Series reflects the Red Flag's brand-new design of Linux server product in 2002-2003, it includes many innovations of Red Flag itself. All new products are all developed in accordance to the strict operating system standards, as well as the open operating system standards and protocols. With large sum of components upgrading (426 packages are updated, 200 packages are substituted) and the user interface is re-integrated. Red Flag Server 4.0 Series possess many new features comparing with Server 3.0 series. From product form to packaging, as well as the certain application positioning, more professional designs are made for the new products. To learn more about major areas of improvements of Server 4 Series visit their Web site:

<http://www.redflag-linux.com/server4/efeatures.html>

You can learn more about Red Flag Linux at their Web site:

<http://www.redflag-linux.com/eindex.html>

1.15.3 Debian Linux

The Debian distribution is a non-commercial distribution whose goal is to effectively compete in the commercial market. This is one reason it is favored by many developers. Debian comes from the names of the creators of Debian, Ian Murdock, and his wife, Debra.

Debian was started in August 1993 as a new distribution that would be made openly, in the spirit of Linux and GNU. Most of the basic operating system tools come from the GNU project, hence the name GNU/Linux. Debian GNU/Linux provides more than a pure operating system.

It is a free distribution and the largest in terms of packages (has over 9,000 packages). The distribution supports multiple languages including French, German, Italian, Japanese, Portuguese, Spanish, Catalan, and Danish. It also

supports several platforms including Intel, alpha, pa-RISC, PowerPC, and sparc64.

What's in the distribution

As mentioned earlier, there are over 8710 packages available in the Debian distribution. You can install the standard applications for graphics, Web browsers, word processors, desktop managers, and so on. There are way too many packages to list here. You can search the contents of these packages on the Debian Web site at:

<http://www.debian.org/distrib/packages>

You can find out more about Debian at:

<http://www.debian.org>

1.16 Going nuts with kernels

The kernel is the central module of an operating system. It is the part of the operating system that loads first, and it remains in main memory. The kernel is responsible for memory management, process and task management, I/O, and disk management. It manages files and is the part of the operating system that interfaces with hardware devices.

The open source portion of Linux that will *always* remain free is the kernel. Hundreds of developers submit their proposed changes to the kernel on an on-going basis. Linus Torvalds and his committee must approve these changes for them to be incorporated into the next kernel version.

Linux kernel versions are divided in two series: developmental (odd series: 2.3.x, 2.5.x) and production (even series: 2.2.x, 2.4.x, 2.6.x). The developer series generates fast moving versions for testing new features, algorithms, and device drivers. Production or stable kernels have a well-defined feature set, a low number of known bugs, and tested drivers. They are released less frequently than the development kernels.

Linux distributions are usually based on chosen stable kernel versions, but not necessarily the latest production version. A different version number of the kernel appears for various platforms (for example, SPARC, Alpha, PowerPC).

You can download the latest Linux kernel from the source:

<http://www.kernel.org>

Another important area of Linux decision making is for workload performance related to scaling, which is in turn determined by the Linux kernel level. Table 1-5 shows the scaling by kernel.

Table 1-5 Linux scaling

Number of CPUs	Linux kernel	Example distributions
One to four	2.4	RHAS 2.1 Red Hat 7.3 SuSE SLES7
One to eight	2.4.18+	UL 1.0 SuSE SLES8 RHAS 2.1
One to sixteen	2.6	First quarter 2004

Do not rely on scaling alone to determine increases in application performance. The way in which the workload performs with the whole system is more important. Adding more processors may not gain as much improvement in performance as achieving a balanced system. Enterprise distributions also have additional advanced functions for scaling.

SMP scaling and Linux are important techniques for clusters and multinode scaling. They include network load balancing, grid computing, Beowulf, OpenMosix, etc. The free per node software licensing of Linux can lend itself to a more cost-effective multi-node solution.

For more information, see the following sites:

- ▶ Load balancing
http://www.lcic.org/load_balancing.html
- ▶ Grid computing
<http://www.gridcomputing.com/>
<http://www.ibm.com/grid/>
- ▶ IBM Redbooks grid portal
<http://publib-b.boulder.ibm.com/redbooks.nsf/portals/OnDemand>
- ▶ Beowulf
<http://www.beowulf.org/>
- ▶ OpenMosix
<http://openmosix.sourceforge.net/>



Open source software

Open source refers to any program whose source code is made available for use or modification as users or other developers see necessary.

Few people know that open source was the business model that software began with! In the 1960s, nobody would buy a computer (a huge investment at that time) that was not immediately ready for some use. Software had to be given away by manufacturers as “a way to sell the hardware faster,” and free of charge for that reason. The source code was distributed so that anybody could change it. At the time, nobody would or could use a computer without having programming skills.

In the 1970s, legislation in many countries decided that giving away software for free was unfair competition toward new companies that were developing competitive software and selling it for a living. Therefore, software had to be billed by law. It was still lawful to give the source for free as an option to anybody who bought a software product, and IBM did just this.

In the late 1970s, centralized support to end-users began to get difficult because they did not know whether they used “vanilla” software (object code from IBM) or software modified by their IT group. Since only very few people ordered the source, but were causing most of the work, IBM had to adopt an object-code-only (OCO) distribution mode. This may have been one of the reasons universities began to switch from VM/CMS to UNIX, because UNIX code was available to play with.

But open source software (OSS) is different in nature. It is usually developed as a public collaboration and made freely available on a machine-independent way. Most of the software is written in C. One of the first tools written by Richard M. Stallman (RMS), founder of the Free Software Foundation, was to write an Open C compiler called the *gcc*.

Why are the manufacturers (including IBM) interested in open source? Because as hardware costs went down by a factor of about 2300 to one in 20 years, software costs per user climbed. The *cost of software* is now the *biggest inhibitor* to the *adoption of hardware* (and for the desktop more than anywhere else). Anything that can lower it helps to unlock the situation.

This chapter introduces you to some key concepts of open source software. It also helps you to understand the relationship of Linux and IBM with open source.

2.1 Creating the code and protecting it

Any discussion of open source software requires an explanation of how software is taken from a more human-readable form (source code) and converted into execution code (binary or object code). It is also important to define the various types of licenses that serve to protect intellectual capital in the form of computer software.

2.1.1 Source code

Most computer programming is done in one or more languages that are abbreviated, compact, and highly structured versions of English. This language has to meet four needs that are difficult to match (therefore, the high number of available languages):

- ▶ It should be translatable into something that a computer can understand (a process called *compiling*).
- ▶ It should be easily readable by a writer or other programmers. The language in its original form is called *source code*. Not many people agree about what is “easily readable”, especially among C, C++, APL, Java, Perl, Python, and Haskell programmers (all computer languages). Example 2-1 shows source code from the Java programming language.
- ▶ It should not be too close to the machine architecture to be easily portable.
- ▶ It should not be too far away from the machine architecture to be efficient.

Note that the two last points are in complete opposition.

Note: Pre-incrementation (++) and post-incrementation (++) in the C language are directly bound to the PDP-11 architecture. However, programmers came to like this feature, and it was kept in the language.

Example 2-1 Java source code example

```
import javax.swing.*;
import com.sun.java.swing.*;
import java.awt.*;
public class HelloSwingApplet extends JApplet {
public HelloSwingApplet() {
    getRootPane().putClientProperty("defeatSystemEventQueueCheck",
                                    Boolean.TRUE);
}
public void init() {
    JLabel label = new JLabel(
        "You are successfully running a Swing applet!");
    label.setHorizontalAlignment(JLabel.CENTER);
    label.setBorder(BorderFactory.createMatteBorder(1,1,2,2,Color.black));
    getContentPane().add(label, BorderLayout.CENTER);
}
}
```

Compiling the code

After the source code is completed, it is then submitted to a *compiler*. A compiler is a software program that reads the special instructions that make up the source code, and creates a file that can be executed by a specific type of processor.

Because computers read their code using binary numbers, this file is said to be a *binary file*. The process of compiling is similar to a cook following a recipe to create a meal. In compiling, many of the structures and labels in the code are removed, just as in the process of preparing a meal, where many of the original food forms are altered beyond simple recognition.

Binary code

The output of the compiler is in a form made up of binary numbers. Example 2-2 shows a section of a compiled program.

Example 2-2 Binary numbers from a computer program

```
00000101 00011111 00001000 00000000 00000000 00000000 11010010
01100001 01100000 01100001 01001000 01011000 01000000 11010010
01000001 11100000 10111000 01101000 11010010 00000011 01000000
11000000 01001000 01011000 11110000 01000000 00001100 00000101
00000100 11000010 00000000 00000000 01000111 11110000 10111000
```

Because the compiler removed many labels, instructions, and forms from the source code in creating the binary code, it can be difficult for a human to determine the nature or structures of a program simply by reading the binary output of the compiler. Most software, especially commercial software, is shipped only in a binary format. The source code is heavily protected from publication. This preserves the methods and technologies of the software creator, who typically invests a significant amount of money in their development.

2.1.2 Software licenses

Most software is published with copyright protection, similar to a book. Unlike a book, however, software typically adds an additional layer of intellectual property protection in the form of a license.

Licenses serve to specify, clarify, and restrict the legally permissible use of a software product. Most licenses require the purchase of rights to use the software. Figure 2-1 breaks down the various software licenses into several categories.

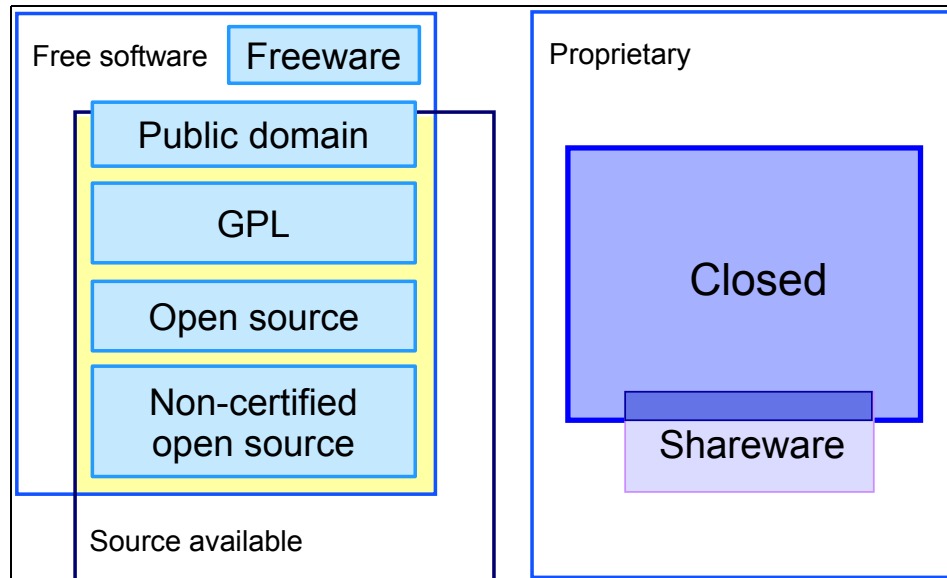


Figure 2-1 Software license categories

Proprietary software includes the following two overlapping subcategories:

- ▶ **Closed software:** This is any software where the source code is not available for review or modifications. However, it can be released as shareware. This is the traditional model of most software.
- ▶ **Shareware software:** This is a common license for many popular programs and applications, including some enterprise-level software. Shareware licenses allow the use of an application, without charge, for a limited time, primarily to allow an end user to evaluate the software. It is common practice to limit the use or functionality of a software program for the duration of the evaluation period. At the end of the evaluation, a user is either prohibited from using the program, or is given the option to purchase special code that allows full supported use of the program.

Free software overlaps heavily with open source software, but not freeware. *Freeware* is any software license that allows you to use the software at no charge and with no expectation for payment in the future. Software issued under a freeware license traditionally does not include the source code. It is usually submitted by individuals or small organizations working on small programs and projects. Let us see how open source software fits into this picture.

2.2 The origin of modern open source

The label *open source* is applied to a philosophy of intellectual property that is the basis of several software licenses (and a general license category formally called open source) and a model of software development. Most proponents think of open source as concept that asserts that the user of a software program should have the right to see and modify the source code that was used to create that program.

The open source movement has its origins in the development and licensing of UNIX systems in the 1970s. AT&T, the creator of UNIX, liberally licensed UNIX systems to U.S. research universities to take advantage of any developments made to the operating system by university workers. Because, starting in 1974, many of these universities were connected to each other through ARPAnet, the precursor to the Internet, individual developers could share improvements and troubleshoot software problems. They could generally participate in a peer-to-peer environment that crossed geographical and organizational lines.

The UNIX system became a commercial-only operating system in 1984. This act ended the liberal code-sharing license that allowed developers to take advantage of each others code improvements. As a result, an attempt was made to continue the tradition of code sharing, most publicly by Richard Stallman, who helped start and still heads the Free Software Foundation (FSF).

2.2.1 The Free Software Foundation

“Outstanding people have one thing in common: an absolute sense of mission.”

- Zig Ziglar

According to the Free Software Foundation, free software allows users to freely run, copy, distribute, study, change, and improve the software. It give users the following types of freedom:

- ▶ **Freedom 0:** Run the program for any purpose (liberty)..
- ▶ **Freedom 1:** Study how the program works and adapt it to your needs. You must have access to the source code for this (equality).
- ▶ **Freedom 2:** Redistribute copies to help others (fraternity)
- ▶ **Freedom 3:** Improve the program and release your improvements so the whole community benefits. Again, you must have access to the source code for this.

You can learn more about free software on the Free Software Foundation Web site at:

<http://www.fsf.org/philosophy/free-sw.html>

The Free Software Foundation believes that denying access to the source code of a software program is morally wrong. Because of this, it created the General Public License (GPL). This license forces the free use (as previously defined) of any software protected by the GPL. A distributor of GPL software can charge to distribute the software code or support services. However, the users of the software can modify or use it in any way without paying the copyright owners for that capability.

GPL-protected software is different than software in the public domain for the simple reason of derivation. Code that is based on public domain software can be copyrighted and closed, while derivatives of GPL software are required (by the GPL) to remain open. The GPL focuses on freedom of use and not freedom from being charged for the software.

2.2.2 The compromise of open source software

GPL was the first license to publicly protect access to software source code. However, most commercial organizations felt that the license did not provide enough protection for their investments in software development. Other organizations were more concerned with controlling the integrity of their

applications, and limiting the publication of derivative works that might fragment or compete with their control of the original program.

As a result, the 1980s saw the creation of many similar-themed, but less restrictive licenses such as the MIT X license or the BSD license. In 1998, when Netscape announced that they were going to release the source code of their Navigator Web browser to the public, a viable alternative to the GPL was created.

The open source initiative

The phrase “open source” was first proposed during a meeting in February of 1998 of well-known Linux proponents including John “maddog” Hall (from Linux International), Sam Ockman (a member of the Silicon Valley Linux User’s Group), and Eric Raymond (author of *The Cathedral and the Bazaar*). Netscape asked Eric Raymond to help with the process of releasing and managing the code. He wanted to use the opportunity to help push the open source idea and development model into the commercial arena.

As a result of that meeting, the open source initiative was created. This not-for-profit group created a formal description of open source called the Open Source Definition (OSD). It also created a certification process for any license that attempted to bill itself as open source.

Note: Most popular open source software licenses are certified as OSD compliant. You can find a list of them on the Web at:

<http://www.opensource.org/licenses/>

The OSD outlines use of the software based on the following criteria:

- ▶ **Free distribution:** You can sell or give away open source software as part of a greater software package that includes other programs from any number of sources. This eliminates the possibility of the original license imposing redistribution fees or royalties.
- ▶ **Source code:** You must distribute open source software with the source code. Or you must make it easy to obtain the source code without an unreasonable reproduction charge. You must present the code in a common and non-misleading format. If source code is not easily accessible, it is more difficult to improve or evolve the code.
- ▶ **Derived works:** Anyone is allowed to create modifications to the original code and distribute those modifications under the same terms as the original software license dictates. This allows the cycle of improvement and development to continue, while preserving the control of the original developer of the code.

- ▶ **Source code integrity:** While a software license cannot prohibit changes to the source code, it can limit how that modified code is distributed. The most common restriction is to specify that any changes must be distributed as separate patches that can be integrated into the original source code (or ignored) at the discretion of the user. Another method of achieving the same aim is to force modified code to carry a different name or version from the original. Both restrictions serve to establish the responsibility for software changes while preserving the brand or product identity.
- ▶ **Personal discrimination:** Open source licenses are not allowed to restrict any individual or group of individuals from using their software. However, they can warn users about legal restrictions or obligations. The restrictions cannot be built into the software itself.
- ▶ **Use discrimination:** An open source license cannot restrict how the software can be used. This means that any industry or endeavor can take advantage of the software.
- ▶ **License distribution:** All rights granted by an open source license must be without provision. That is, no further agreements (such as a non-disclosure agreement) or licenses can be required for consideration in compliance with the original license.
- ▶ **Non-product specific:** The permissions granted as part of an open source license cannot depend on the code being distributed as part of a package or group of software. Modifications to a specific licensed product can be distributed without including the entire package (or additional products) that made up the original distribution.
- ▶ **Non-restrictive:** An open source license cannot make demands of other software that is distributed with the original program. For example, a license cannot dictate that all programs distributed with it (as part of a package or media, such as a CD) must also be open source. This protects commercial interest by allowing the bundling of commercial software with open source software.
- ▶ **Technology neutral:** This point aims to remove the requirement that a user performs a specific action (other than installation and use) to accept the open source license. This allows you to distribute the software through multiple platform and media types without requiring a graphical user interface (GUI) window option, for example, to accept the license.

You can learn more about the OSD on the Web at:

<http://opensource.org/docs/definition.html>

These ten points were an attempt to bridge the gap between an apparent philosophical dispute between two diametrically opposed licensing schemes, the forced freedom of the GPL on one side, and the complete control and secrecy of

most typical closed source software licenses on the other. The core issue is about licenses. It is also about which software development methodology and process is superior, closed and centralized or open and distributed. The implication is about revenue of a finished product and control of the development process itself.

This is the genius of the open source movement. It is a compromise of both ends of the development and licensing spectrum. It allows the original creator of a program to have some control over how the software develops and is presented to the public. At the same time, individual developers and users retain the ability to manipulate and share improvements of the source code for their own needs.

These theoretical arguments originally were not enough to convince businesses to change established processes and culture, and adopt a more decentralized and open developmental model. The success of Linux gave organizations an effective example of how the compromise between distributed development and centralized control could (and did) function successfully. The success of open source is an example of what is commonly called *percolation*.

A new production mode?

Economists like to analyze history using what they call *production modes*. Going from hunt to discovery was a change of production mode and the base of today's wealth.

The OSS model is seen more and more as a *new production mode* that can shape the future. Previously, such mode of cumulative change existed for everything related to scientific research. The Internet simply allowed for the generalization of new (technical) areas and its quick acceleration at the speed of light.

With OSS, for the first time, people who do not know one another work on common technical projects for reasons other than fear or the need for money.

2.2.3 Why the OSS model works so well

Two factors are at work here: the human hierarchy of needs as described by Maslow and strict self-interest considerations.

Human hierarchy of needs

According to Maslow, human needs are made of five priority layers, each of them becoming unimportant as soon as the corresponding needs are met. These needs are:

- ▶ Basic metabolic needs
- ▶ Security

- ▶ Recognition by peers
- ▶ Self-appraisal
- ▶ Being useful to the rest of the world

Where the first two are fully satisfied, the last three may take the lead. Open source software developers are in these three layers and happy to be there.

Self-interest considerations

There are two ways to increase your purchasing power. One is to get more money. The other is to start buying so that things becomes cheaper.

Let us suppose that you devote 40 hours of your time each year to add some improvement to an open source product and put it in the common pot. If 10 000 people do the same and you are interested only in 1% of what they do, you receive 100 times more than what you give after just one year. This is a return on investment (ROI) of 1000% per year, which is much more than any financial placement can offer.

But even Linus Torvalds himself did not expect such a *snowball* effect.

2.3 Linux, Linux, and open source

Linus Torvalds began working in 1991 to create a better terminal emulator for his PC. Within a few months of working in his spare time, he created a small but workable clone of the UNIX system kernel. He published this (with the source code) on an Internet news group that was dedicated to operating system kernels. Several news group participants immediately offered ideas (and actual code) as improvements. This began a process that continues today of individuals and organizations providing ideas, new code, and patches that are evaluated by Linus (and a select group of volunteers) for inclusion into the kernel.

You can find a well-researched overview of the Linux development process on the Web at:

http://firstmonday.org/issues/issue5_3/kuwabara/index.html

2.3.1 The GNU utilities

A crude operating system by itself was not enough to convince more than merely interested to take it seriously. Nor was Linux the first project that relied on an open source development model. That title belongs to the Free Software Foundation. Under Richard Stallman's direction, it had created the GPL and released numerous utilities under the license as part of the GNU project. These included the GNU C Compiler Collection (gcc), which currently contains front

ends for C, C++, Objective-C, Fortran, Java, and Ada, as well as language libraries (libstdc++, libgjc, etc.). It is a high-performance, multi-platform compiler that remains the standard by which all other compilers are judged.

You can locate the GNU C Compiler Collection on the Web at:

<http://gcc.gnu.org>

Because Torvalds decided to release Linux under the GPL, it was natural to bundle the GNU utilities with it. In fact, the Free Software Foundation argues that because so much of Linux was built on, and distributed with, the GNU utilities, Linux should be called *GNU/Linux*. While little argument exists against the case made by the Free Software Foundation, few have adopted this label, preferring the simpler name of *Linux*.

2.3.2 The problem of complexity

As the development of Linux continued, it was inevitable that significant complexity would emerge. This was especially true as the kernel grew to over one million lines of code. This code was mostly maintained by part-time volunteers (estimated to exceed a total of 40,000 in number) worldwide, who communicated with each other almost exclusively through e-mail.

This complexity is managed by Linux developers using a version control program called the Concurrent Versions System (CVS). CVS works by recording all the changes that occurred in the source code of an application. As a result, bug or error detection is simplified because it is possible to track the exact evolution of events that led to a particular software error.

You can find the CVS on the Web at:

<http://www.cvshome.org>

CVS also works to enable software coordination. Each software developer has their own unique work space. This isolates them from the work of others until their code is ready to be merged, a process that is aided by CVS.

But CVS is only part of the answer to the problems created in distributed development. For example, coordination and communication between developers is still necessary to determine when and how code should merge or branch, whether multiple changes within a file or across multiple files will functionally or logically conflict, and how to resolve those conflicts.

CVS is a tool to aid management and control of the development process, but it does not create or manage that process by itself. How then does the chaotic model ascribed to open source software allow for something so complex as an

operating system, let alone an operating system that is stable enough to be considered for deployment in critical enterprise applications?

2.3.3 The benevolent dictator

The development of Linux is mostly distributed and decentralized. However, Linus Torvalds himself still exerts authority and ultimately maintains coordination, along with an informal group of trusted, knowledgeable volunteers colloquially called the “Inner Circle.”

Linus tends to focus on developing the beta version of each kernel version. After he considers a version stable enough for general release, he turns over the responsibility of maintaining that stable version to Alan Cox, another well-known programmer. While Linus turns his attention to the next beta kernel, Alan is left with wide discretion about how the current kernel is maintained and updated.

This same pattern is seen with other members of the Inner Circle. That is, an individual is in charge of development and maintenance of an area they originated or took over from another original developer.

Because of the sheer complexity and quantity of the kernel code, its components, and the various patches and feature requests, Linus has encouraged a system where any submitted code is first sent to the individual who oversees the area in question. As a result, Linus’ principles and goals in regard to the development of Linux are still maintained, while significant authority is delegated to the individuals who oversee specific subareas.

In this fashion, Linux follows more of an open source model of development. It combines the innovation that comes from forced free code imposed by the GPL, yet retains the vision of the kernel’s original author through a delegation of authority.

2.4 IBM and open source

IBM is committed to open source as both a license and a development model for several reasons:

- ▶ IBM clients and partners have requested open source software (including Linux) support for all IBM platforms, products, and solutions.
- ▶ Open source software, with its wide distribution and use, typically becomes an industry standard.
- ▶ Innovation within the open source community typically occurs at a higher rate and volume than in closed-source communities.

2.4.1 IBM and open source licenses

The first IBM open source license was the IBM Public License (IPL), introduced in 1999 and certified by the OSI in the same year. As part of the Eclipse release in 2001, IBM updated the IPL and renamed it the Common Public License (CPL). The OSI certified the CPL in May 2001 as an open source license.

For more information about the CPL, see:

<http://www06.ibm.com/developerworks/library/os-cplfaq.html>

2.4.2 IBM and open source software

IBM is a major contributor to the open source community, with over 250 developers worldwide working full time on several initiatives. Such initiatives include:

- ▶ **Linux Technology Center (LTC):** The LTC works to enable the enterprise capabilities of Linux through development and contribution of technology, utilities, tools, and code.

You can learn more about the LTC on the Web at:

<http://oss.software.ibm.com/linux>

- ▶ **Open source software support:** IBM is an active supporter of projects and software including:
 - Open Source Cluster Application Resources (OSCAR)
 - Open Source Developer Lab (OSDL)
 - GNOME
 - KDE
 - Open Source Initiative (OSI)
 - Free Standards Group
 - USENIX
 - Linux high availability
 - OpenLDAP
 - USB
 - PCI hot plug
 - Advanced Power Management (APM)
 - PPC-32
 - PPC-64
 - Stream Control Transmission Protocol
 - Free Standards Group (Linux Standard Base)
 - Samba

Learn more about the involvement of IBM in open source software on the Web at:

<http://www.ibm.com/developerworks/oss/>

- **Eclipse:** This is an open source software integrated development environment (IDE) sponsored by IBM that acts as a development infrastructure platform. Through extensions (called *plug-ins*) to the base IDE, Eclipse can provide nearly any type of capability, including code versioning functionality and even syntax highlighting. IBM WebSphere Studio Workbench is a commercial product that is built on Eclipse.

You can find Eclipse on the Web at:

<http://www.eclipse.org>



Linux, UNIX, and Windows systems

The majority of enterprise Linux implementations are found in corporate Internets or in the Internet infrastructure. Other notable uses include file and print server deployment, application development, messaging server deployment, data warehousing, embedded applications, thin client applications, and massively parallel scientific applications. Linux is becoming robust enough to find its way into engineering software development, the healthcare industry, the retail industry, and financial institutions.

Wall Street has adopted Linux in a big way. Major companies have been in the press, sporting their new implementation projects with Linux. These are not just departmental solutions, but large company-wide solutions indicating potential savings of 30 to 1. These results will certainly attract the attention of Chief Financial Officers (CFOs) and Chief Information Officers (CIOs) looking to improve their business with leading edge solutions and major spending reductions.

Linux is also taking hold in the public and government sectors as the IT industry tries reduce expenditures. How does it compare to the traditional UNIX systems, vendors, or Windows? Does Linux provide a real advantage? This chapter provides a high level overview of some of the key characteristics of the Linux, UNIX, and Microsoft Windows operating environments.

3.1 Cost of ownership

This section discusses TCO as it relates to the planning and implementation of Linux systems. It also contrasts TCO with return on investment (ROI).

3.1.1 Total cost of ownership

Vendors, consultants, and market research firms have done numerous total cost of ownership studies. As shown in Figure 3-1, most research seems to agree that there is a TCO advantage for Linux in long-term deployment and operation. Most research illustrates some type of savings with the deployment and operation of a Linux environment, although the amount of savings can differ greatly.




Figure 3-1 Typical Linux TCO study showing cost over time

Many factors can skew a simple TCO study. Such factors include:

- ▶ A comparison of the Linux and non-Linux operating environment, also which Windows operating environment, UNIX operating environment, and Linux distribution
- ▶ Whether a list price or deeply discounted price was used.
- ▶ Whether new hardware was configured into the TCO model.
- ▶ Whether a desktop deployment, enterprise deployment, or SMB deployment was used in the TCO model.

- ▶ The type of workload used in the TCO model, perhaps including collaborative, Web server, file server, or database server
- ▶ The software and applications used in the configuration
- ▶ Assumptions about administrative capabilities of the chosen configuration and the number of administrators needed
- ▶ The time frame the TCO model covers
- ▶ Whether support and service contracts were configured into the model
- ▶ Whether consideration was given to training a user when transitioning to a different operating environment and the short term loss in productivity
- ▶ The sponsor of the study (to some degree)

The cost savings of Linux may not always be as clear and simple as many open source advocates have proclaimed. You must factor in any hidden costs into the TCO model. Such hidden costs may include the skill level of the IT staff, the complexity of applications, and the cost to migrate from another environment.

The purchase price of hardware and software alone provides only a small indication of the total lifecycle cost for an IT solution. The TCO life cycle includes all staffing costs associated with acquiring, maintaining, and removing an IT asset from an organization.

There are numerous models to categorize the total cost of computing. You can break down the total cost of computing model into the following categories:

- ▶ Hardware
- ▶ Software
- ▶ Staffing
 - Development
 - Support

You can also break down the total cost of computing model into the following categories:

- ▶ **Buy:** Includes making the lease or purchase decision
- ▶ **Use:** Setup, training, use, and technical support
- ▶ **Dispose:** Removing the asset, transfer data, and applications

Most TCO models do not include the cost of external communication, media, power consumption, floor space, etc. since they are somewhat similar. The key to any TCO model is the categories to be measured, and the quality of the measurements.

You cannot calculate other key factors normally into a TCO model. Such factors may include complexity, scalability, security, manageability, frequency of

software upgrades, availability of applications, and reliability. Discounting can also play a major role in the TCO of a solution. In many instances when a complete solution is purchased from one vendor, deep discounting is applied.

There is also a productivity paradox with selecting an operating environment that is not immediately measurable in any of the TCO models. For example, will the business become more productive or make a better return on investment if the business uses a Linux, UNIX, or a Windows operating environment? Does the decision to move to a Linux, UNIX, or Windows operating environment affect the basic business structure?

3.1.2 Operating environment and product TCO perspective

While it can be difficult to qualitatively analyze the TCO of a total solution that has many intangible aspects, you can analyze some of the tangible costs of the hardware and software. Table 3-1 identifies costs that you can analyze in a Windows versus Linux operating environment solution.

Table 3-1 Cost analysis comparison in a Windows versus Linux environment

Function, product	Windows solution	Cost	Linux solution	Cost
Base operating environment	Windows 2003 Server	\$\$\$\$\$	Red Hat, SuSE, Turbolinux boxed set	\$\$\$
Web server	IIS	\$	Apache	0
Antivirus software	Norton Antivirus	\$	Squid, Virulator	0
Database	SQL Server	\$\$\$\$\$	PostgreSQL	0
Mail server	Exchange Server	\$\$\$\$\$	Sendmail, Postfix	0
Office suite	MS Office	\$\$\$	StarOffice, OpenOffice	\$
Development environment	VisualStudio	\$\$\$	KDE	0
Desktop client	Windows XP Professional	\$\$\$\$\$	Red Hat, SuSE, Turbolinux	\$
Note: 0 = zero cost; \$ = low cost; \$\$\$ = medium cost; \$\$\$\$\$ = high cost				

Function, product	Windows solution	Cost	Linux solution	Cost
System management	MOM, SMS, AC	\$\$\$	Webmin	0
Note: 0 = zero cost; \$ = low cost; \$\$\$ = medium cost; \$\$\$\$ = high cost				

3.1.3 Findings

The following list outlines general, key findings with most TCO models that compare and contrast the Linux, UNIX, and Windows operating environments:

- ▶ For many businesses, the Linux operating environment has emerged as a viable alternative for enterprise computing. Using the Linux operating environment has the potential to lower IT costs. This is especially true for businesses that have the right mix of staff, skills, commodity hardware, software, and overall computing environment.
- ▶ To realize the full potential of the Linux operating environment, companies may need to pilot a Linux project, and then build a vision or strategy for a longer term deployment.
- ▶ The cost of Linux and Windows operating environments on Intel platforms is highly competitive and low due to the acquisition cost of Intel architecture server hardware.
- ▶ Linux can be deployed with no or low licensing fees, particularly if Linux is deployed without the support services that normally accompany an enterprise-level product. This fact alone can account for a lower TCO.
- ▶ TCO may depend on the configuration and workload. However, with some key workloads, such as Web applications, collaborative applications, firewall, file server, or print server applications, Linux is likely to have a lower TCO.
- ▶ Most companies, whether large or small, operate in at least a dual operating environment.
- ▶ Staffing and skills are a key consideration.
- ▶ Linux can be used as a tool for negotiating the price of proprietary solutions in other operating environments.

The IBM *@server* offers the possibility for server consolidation. As explained in the following chapters, these servers have the capability of virtualization, from hundreds of virtual Linux servers on zSeries to tens on other IBM *@server* platforms. This provides additional TCO savings.

3.1.4 TCO versus ROI

While TCO is an important consideration, ROI is equally significant in the customer decision-making process. CFOs take a good look at the ROI they can expect to realize with new servers. While it may be useful to consider the overall total cost of ownership, today customers are interested in saving money in a tough economic environment.

ROI can be more effective than TCO as a measure of business value, especially with new, emerging technologies. For this reason, we expect to see a movement of the community toward a more comprehensive ROI justification, where TCO is just a portion of the overall analysis.

It is coming to be viewed that ROI analysis is a higher quality approach for making business decisions since it is more inclusive of the risks and returns of making new IT investments. This is what CFOs are looking for, and why Linux is becoming more accepted each day.

Numerous case studies have shown Linux to deliver outstanding ROI value. These benefits may derive from much greater deployment flexibility, much greater openness, much higher performance and functional benefits (derived from the prevalent deployment architectures described earlier), and universally accepted better reliability and security. Such emerging technologies as Linux typically excel in ROI analyses, which highlight the benefits of the new technology, one of which is cost savings.

Linux today contributes to business innovation through the following typical high-value deployment methods that provide TCO and ROI advantages:

- ▶ **Using appliance servers** to provide fixed, single-function solutions simplifies use and lowers support, operation, and maintenance costs. Single-function solutions are often deemed preferable for infrastructure workloads over adding software to a general-purpose server.
- ▶ **Consolidation of distributed workloads** onto centralized servers dramatically reduces the overall administration costs and delivers improved resource utilization.
- ▶ **Clustering Linux servers** for high-parallel or repetitive tasks simplifies administration (since tasks can be scripted and repeated across multiple servers). This greatly enhances performance. Customers are deploying Linux widely, often in high availability configurations, which are more suitable, reliable, and secure with Linux.



zSeries and Linux

This chapter provides a brief overview of the Linux solution that is available on the zSeries server. It includes information about the following topics:

- ▶ The various modes that Linux runs on the zSeries server
- ▶ The benefits of running Linux as a z/VM guest
- ▶ Networking options that are available for z/VM Linux guests
- ▶ Storage options that are available for Linux on zSeries
- ▶ Consolidation scenarios using Linux on zSeries
- ▶ Software solutions for Linux on zSeries
- ▶ Special IBM offerings for zSeries and Linux

4.1 Introduction: Linux on zSeries

Reduced IT budgets, increasing security threats, proliferating islands of information cut off from the people who need it, and the pressure to satisfy requirements quicker than ever: these are just a few of the obstacles facing organizations of all sizes as they strive to respond to customer needs and market opportunities.

Enter Linux for IBM @server zSeries, which can help transform your business into a true on demand operating environment. Built on open standards, Linux running on zSeries servers gives you freedom of choice in middleware. zSeries solutions incorporate workload management technology that can enable your Linux applications to share computing resources and to allocate them dynamically as needed. And with advanced virtualization capabilities, zSeries servers running Linux can empower you to accomplish more with fewer resources.

The TCO advantage

Linux on zSeries can help lower your total cost of ownership (TCO). Whether you are a small, medium, or large enterprise, over time Linux on zSeries can be a boon to your bottom line in a variety of ways:

Possible savings on hardware

Is your server farm growing out of control? Linux running on zSeries can give you the ability to rein in server sprawl through consolidation. By exploiting the virtualization capabilities of zSeries, you can consolidate hundreds of individual Linux servers onto a single server, and deploy these virtual servers in a matter of minutes. Fewer servers can also mean additional savings in physical space, and heating, air conditioning, and electricity costs.

Possible savings on software

Because Linux is an open operating system, running your applications on a zSeries server with Linux helps save you from getting locked in to costly licenses for proprietary software. If you have already put money into certain software to run your core business applications, you may be able to implement that software on Linux on zSeries, which allows you to run diverse applications and middleware alongside Linux in multiple dynamic partitions. Consider also that Linux application software and middleware is usually priced on a per-engine basis. Since a single zSeries server can run multiple Linux applications on a single engine, your licensing costs can be reduced by consolidating distributed servers onto a zSeries server.

Possible savings on staffing

Consolidating multiple servers onto a single zSeries server running multiple virtual Linux servers can mean less labor required for system management and maintenance. The centralized system management and autonomic computing capabilities of Linux on zSeries can also help cut down on the errors and workload-balancing tasks that otherwise can eat up countless IT staff hours.

Possible savings on business continuity

Linux on zSeries can give you the ability to scale up and out *on demand* to meet spikes in server activity, helping to minimize costly transaction delays and potentially devastating system crashes. With a suite of built-in self-healing, self-managing, and self-optimizing features, zSeries servers can rapidly respond to, or even anticipate, threats to system health, helping to prevent budget-sapping downtime.

4.1.1 Advantages on Linux on zSeries

Linux on zSeries is all about helping to remove complexity from your IT infrastructure: reducing server sprawl, keeping a lid on software licensing fees, and minimizing the need for human intervention in managing and maintaining your servers. Here are the key capabilities of Linux on zSeries that make this possible.

Virtualization

Virtualization technology is built into every server in the zSeries line. It allows you to create and manage discrete virtual processors and communications, storage, and I/O devices in a single zSeries server. These resources can be dynamically shared or re-configured, helping to simplify the demands placed upon you to plan, purchase, and install new hardware to handle new workloads. Also remarkable is the speed with which zSeries virtualization technology allows you to create and deploy virtual Linux servers — in minutes rather than days. This can help accelerate your project schedules, and ultimately, your time to market.

Business Integration

Based on open standards, Linux running on zSeries servers is designed to provide high-speed connections among heterogeneous applications in the same server. These high-speed connections are called HiperSockets™, which can be described as an in-storage implementation of a TCP/IP network. By reducing the need for an external network, HiperSockets helps you reduce costs, and complexity, and since all of the communication takes place within the zSeries server, security is improved as well. zSeries servers give you the choice of running multiple Linux, z/OS®, z/OSe, z/VM, TPF and VSE/ESA™ systems at the same time, along with diverse middleware components. The result is a highly integrated environment designed to be free of barriers to information sharing,

allowing your business to become more responsive, flexible, and efficient, which can help you to get a unified view of your customers and competitors, launch new products faster, and process transactions more smoothly.

System management

Advanced zSeries system management tools are designed to provide the proper amount of resource to the most important applications when needed. The VM Resource Manager lets users assign performance goals to help ensure that your on demand business meets demands for service. And, Linux together installed with z/OS on zSeries allows you to take advantage of the z/OS Workload Manager, which can manage the logical partition that Linux is running in, and in concert with the Intelligent Resource Manager to dynamically balance your computing resources to your workloads, helping you to be prepared for spikes in demand.

Big jobs get the muscle they need, and small jobs only what they need, keeping utilization and availability high. Virtualization software gives you a single point of system management control, helping to minimize errors, and by extension, minimizing effort and costs. And innovative self-configuring, self-healing, self-optimizing, and self-protecting features built into every zSeries server can keep your system running at peak performance with a minimum of human intervention.

Resiliency and security

zSeries provides a highly secure Linux platform. zSeries servers today are the result of over 40 years of innovation. Designed to be rugged, resilient, highly available, and highly secure, zSeries servers are often entrusted with the most important, mission critical applications and data that drives a company. Autonomic attributes are designed to allow your machine to be there when you need it, and you can add compute power to a zSeries server to meet unexpected demand when you need it. Also, virtual Linux servers on zSeries are isolated, so you can run multiple, heterogeneous workloads simultaneously. zSeries Linux solutions can also incorporate several industry-leading security features such as enhanced cryptography support designed to allow e-business transactions to remain private and secure.

4.1.2 Ways to run Linux on zSeries

There are three ways to run Linux on the zSeries server. Each way is explained in the following sections.

Basic mode

Linux can run on the entire machine, without any other operating system. In basic mode, you can use a single Linux image on the zSeries “bare metal.” Although it is possible to run Linux on zSeries servers in basic mode, it is not recommended since Linux is typically used to support a single application, it would not be economically viable. This option is not supported with zSeries 990 models. Typically, Linux on zSeries servers runs in one or more logical partition, and as a guest of z/VM.

Logical partition

You can logically partition the zSeries hardware into a maximum of 15 separate LPARs. A single zSeries, for example, can host z/OS applications in one partition, VM and VSE applications in others, and Linux applications in additional partitions.

With zSeries 990 Exploitation support enabled, you will be able to use up to 30 LPARs. This feature is now available for z/VM 4.4. A Linux open source code-drop is available on Developer Works and will be incorporated by distributors with next releases.

For more information about new hardware features available with latest kernel upgrades, see developerWorks navigation for Linux for zSeries at:

<http://oss.software.ibm.com/developerworks/opensource/linux390/index.shtml>

z/VM

A customer can run Linux as a virtual machine using z/VM as a hypervisor. z/VM in concert with this functions in the zSeries hardware, provides virtualization of CPU processors, input/output (I/O) subsystems, and memory. A customer running z/VM can run hundreds of Linux systems on a single zSeries. With z/VM, for instance, a customer can offer a complete Linux server environment to each application developer, and host the production system all on the same zSeries server. We discuss z/VM and Linux in more detail later in this chapter.

The decision on how to run Linux on zSeries is typically made based on the expected workload:

- ▶ Linux workloads with a large memory footprint and requiring extensive processing power may be best suited for running in an LPAR.
- ▶ Few (if any) workloads require running Linux in basic mode.
- ▶ Running Linux as a z/VM guest offers the maximum flexibility. The number of Linux guests running under z/VM can be adjusted to meet workload

requirements while maximizing usage of the zSeries hardware. Most customers choose this option to maximize their investment.

4.1.3 31-bit and 64-bit options

Linux solutions are available for the 31-bit and 64-bit environments. The option availability depends on the zSeries model and the Linux distribution.

Note: The choice between the 31-bit and 64-bit addressing environments can depend on the application or middleware you are planning to use. Some software does not support 64-bit addressing mode and may work not properly.

Linux for S/390 (31-bit)

Linux for S/390 is a 31-bit version of Linux. It is currently available for S/390 Parallel Enterprise Server G5 and G6 and the Multiprise® 3000 processors. It also runs on zSeries models in 31-bit mode. Because of the 31-bit limitation, the addressable main storage is limited to 2 GB.

Linux on zSeries (64-bit)

Linux on zSeries supports the new 64-bit architecture on zSeries processors. This is all zSeries 800, 900, and 990 models. 64-bit support eliminates the 31-bit storage limitation of 2 GB.

4.1.4 Integrated facility for Linux

Adding capacity to an existing mainframe to run Linux workload or workloads can be done cost-effectively using IFL engines. The addition of IFL engines does not change a server's model designation. Also, the cost of existing non-Linux software is not affected by the presence of IFL engines.

Figure 4-1 shows a zSeries 900 Model 1C6 configured with three IFL engines. The engines are allocated from the set of available spare CPs on a zSeries server (including the G5/G6). A zSeries 900 Model 1C6 (a six-way, standard engine configuration) has 10 additional processor engines available for use as IFLs, standard processors, or ICFs.

In Figure 4-1, the customer purchased and allocated three IFL engines (IFL0, IFL1, IFL2). The cost of this z/VM environment is one-time charge (based on a three-engine pricing model) in addition to an annual service and support fee.

The presence of the IFL engines does not change a server's model designation. This is still a zSeries 900 Model 1C6. The software fees for the products running

on the standard engines (CP0 to CP5) are not affected by the addition of the IFL engines.

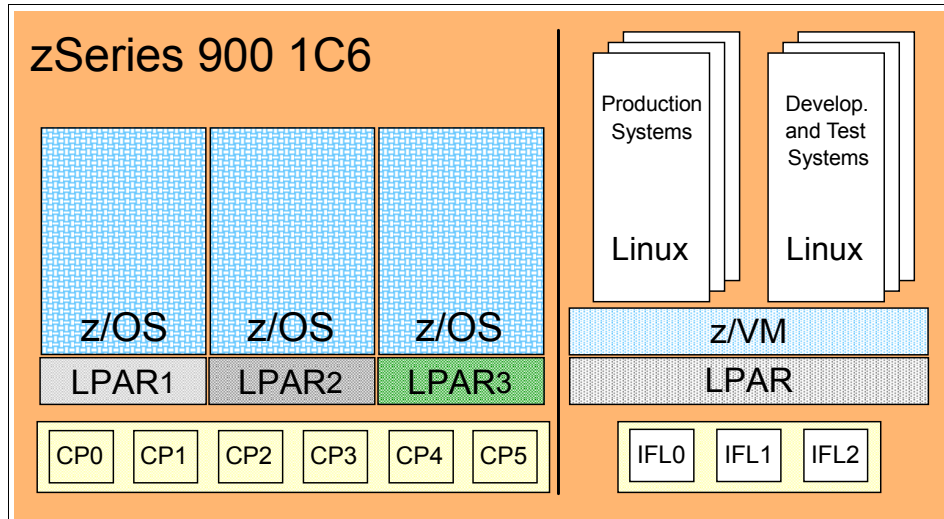


Figure 4-1 zSeries 900 Model 1C6 with three IFL processors and z/VM V4

4.1.5 IPLA pricing

International Product License Agreement (IPLA) is a contract that applies to a type of software from IBM, where the customer pays an up-front, one-time license charge, and an optional software maintenance fee called *Subscription and Support* (S&S). S&S is an annual charge that provides customers access to IBM technical support and enables customers to obtain no charge version upgrades for an IPLA product.

For more information, see:

<http://www.ibm.com/servers/eserver/zseries/swprice/zipla/>

4.2 S/390 and zSeries architecture and features

The following sections describe the S/390 and zSeries in more detail.

4.2.1 Architecture

The heart of zSeries and S/390 hardware is the multichip module (MCM) that contains up to 20 processing units (PU). On other platforms, these are commonly referred to as central processing units (CPUs) or engines.

The core computing elements of any computer are the processing units. These contain the z/Architecture™ logic and functions that implement the architectural extensions to S/390 ESA architecture. As implemented in zSeries Models 800 and 900, each PU actually has dual internal instruction processors.

Instructions are executed by both internal processors in parallel, and the results are compared. If the results do not match, an instruction retry process is performed. This is all done automatically by the PU, and is not visible to the operating system. The normal result of the dual processors in a PU is the execution of a single instruction stream.

We normally refer to a PU as a single processor, and ignore the fact that there are really two parallel processors inside each PU. PUs are used in one of five ways:

- ▶ **Central Processor (CP) PU:** Used by the operating system for executing customer work
- ▶ **System Assist Processor (SAP) PU:** A System Assist Processor is a PU that runs the channel subsystem Licensed Internal Code to control input/output (I/O) operations. One of the SAPs in a configuration is assigned as a master SAP, and is used for communications between the multichip module and the support element.

In logical partition (LPAR) mode, all SAPs perform I/O operations for all LPARs. The zSeries 900 12-PU MCM models have two SAPs as standard. The 20 PU MCM models have three SAPs as standard configurations.

- ▶ **Spare PU:** This is a PU that is not enabled for any purpose. The system uses this PU to replace a failing processor, if necessary. If all four PUs (plus an SAP) are enabled, then there are no spare PUs. Spare PUs may be used for various upgrade options such as Capacity Backup (CBU) and Capacity Upgrade on Demand (CUoD).
- ▶ **Integrated Facility for Linux (IFL) PU:** This type of processor is restricted to running Linux with or without /VM.

- **Integrated Coupling Facility (ICF) PU:** This type of processor is used to run the Coupling Facility function for use in a Parallel Sysplex® environment.

An IFL is a processor reserved for Linux (or Linux under VM). The significance is that it cannot be used to run other operating systems, and its existence is not reflected in the system model number, MIPS rating, or other power ratings. The system model, MIPS, or other power rating method has significant implications for software costs. Adding an IFL does not affect these costs, permitting the use of Linux without impacting other software costs.

Figure 4-2 shows how IFLs may be allocated to LPARs.




Figure 4-2 The relationship between IFLs and LPARs

Important: IFLs must be dedicated to LPARs running Linux or z/VM. They can be shared by multiple LPARs, provided that the LPARs are running Linux or z/VM. Also, an LPAR cannot contain a mix of IFLs and CPs.

An ICF is used only to run the Coupling Facility licensed code. It cannot be used to run normal operating systems. It is similar to an IFL in that its existence does not change the system model number (or MIPS rating) nor impact software costs for system.

IFLs and ICFs require the use of LPARs. If neither of these are used, a zSeries 800 system can run in basic mode (no LPARs). Note, however, that use of the new z/OS operating system package requires the use of LPARs.

This chapter discusses zSeries and Linux. Here are few useful links to additional information:

- VM and Linux for zSeries Resources
<http://www.vm.ibm.com/linux/>
- Linux for IBM @server zSeries
<http://www.ibm.com/servers/eserver/zseries/os/linux/>
- zSeries mainframe servers
<http://www.ibm.com/servers/eserver/zseries/>

Also, several service offerings are available. Table 4-1 lists the platform-based information for services that are available for zSeries servers.

Table 4-1 zSeries Linux information matrix

Information	Location on the Web
Linux for zSeries applications	http://www.ibm.com/servers/eserver/zseries/solutions/s390da/linuxproduct.html
Linux for zSeries operation system download	http://www.ibm.com/servers/eserver/zseries/os/linux/dist.html
Linux for zSeries Development tools	http://www.ibm.com/servers/eserver/zseries/os/linux/ldt/
Linux for zSeries Library	http://www.ibm.com/servers/eserver/zseries/os/linux/library/index.html
Linux for zSeries Education	http://www.ibm.com/servers/eserver/zseries/os/linux/ed.html
Linux for zSeries Customer Stories	http://www.ibm.com/servers/eserver/zseries/os/linux/stories.html

4.2.2 Partitioning and virtualization

Partitioning and virtualization are complementary technologies in that they are most effective when combined – the two concepts working together offer a unbeatable solution – but it is very important to understand the distinction between the two. Simply put, partitioning provides the ability to divide physical system resources into a number of distinct, isolated regions that operate independently from each other – in general, there is a one-to-one relationship between a physical resource and the region it is assigned to, creating the equivalent of a “box within a box” – all the physical pieces behave and perform exactly as they do if partitioning was not present.

Virtualization takes this concept one step further in that it provides the ability to simulate the availability of hardware that may not be present in sufficient amount (or even not at all!) for there to be a one-to-one relationship between physical resources and virtual machines using the resources. Virtualization uses the available physical resources as a shared pool to emulate the missing physical resources, and is often capable of very fine control over how and to what extent a physical resource is used by a specific virtual machine.

Partitioning

As stated above, partitioning physically divides system resources, but the way to do this differs on the various hardware platforms and can be split into the following categories:

- ▶ Hardware partitioning
- ▶ Logical partitioning (LPAR)
- ▶ Software partitioning

Each offers a specific set of advantages and disadvantages.

Hardware partitioning

Hardware partitioning provides physically isolated environments, thus reducing the effect of an electronic problem to a minimum. In order to actually perform work but still maintain isolation between partitions, hardware partitions are very closely tied to the underlying, actual hardware layout, and must be implemented using only vendor-certified components to ensure correct operation.

Logical partitioning (LPAR)

Logical partitioning provides isolated environments handled by microcode, in most cases without modifications to the operating systems running in the partition (although the operating system may be made aware of additional capabilities if LPAR operation is detected). The granularity of resource management in logical partitions is more precise than that of hardware partitioning, as there is a layer of abstraction between the operating system and the physical configuration of the hardware that does not mandate an architectural tie to the underlying hardware configuration.

Software partitioning

Software partitioning provides a further more granular subdivision and larger flexibility of resources already partitioned through hardware or logical partitioning. This is usually done by implementing a “supervisor” application controlling the hardware, and adding substantial code to the operating system and application running in the partition to make them aware of the partitioning, and have them cooperate with the supervisor in managing the available resources. These partitions are often implemented as “performance groups” to allow different scheduling priorities to be applied to different resource groups.

Virtualization

Virtualization is provided by software (and often, hardware assists) that targets management of actual physical resources to emulate a set of logical resources that can be assigned to one or more virtual systems. There does not need to be an actual correspondence between physical and logical resources for virtualization to be used. In most cases, efficient operation demands that the

emulated resources be similar to the ones present in the physical machine, but this is not required.

In a virtual machine, the operating system (and the applications and users of that operating system) is lead to believe it is running on the actual hardware. No modifications are required to the operating system to operate in a virtual machine; the mapping of these virtual resources to real resources is handled by the virtualization software.

4.2.3 Networking

The new zSeries brand offer many solutions to improve your network connectivity. A key aspect about networking is the new feature implemented on zSeries: *Hipersockets Network*.

HiperSockets technology is available for zSeries servers. HiperSockets is an integrated any-to-any virtual TCP/IP network that provide inter connectivity between multiple LPAR or virtualized images under z/VM. With HiperSockets, you can connect any server running on the same zSeries, and improve response time due to low latency.

HiperSockets was invented to provide highly available, high-speed network connections through a memory bus. It can be considered for an application or infrastructure scenarios, where the network is stressed by a large quantity of data exchanged between servers running on the same zSeries server.

Figure 4-3 shows a possible implementation of HiperSockets over an applications server scenario. When you use a Linux for zSeries technology to produce an application environment, you usually plan to have a three-layer architecture composed by:

- ▶ Front layer: HTTP server
- ▶ Middle layer: Application server
- ▶ Back-end layer: Database server

With the proposed architecture, the environment can be planned with:

- ▶ IBM HTTP Server (powered by Apache) under Linux for zSeries
- ▶ IBM WebSphere Application Server under Linux for zSeries
- ▶ DB2 server under z/OS

If a session is started by a browser over the network, a communication between the three layers is started. To avoid network traffic, implementing a HiperSockets solution can help you. With benefits provided by HiperSockets, your network traffic between the layers is exchanged in the zSeries server at memory speeds.




Figure 4-3 The application server scenario with HyperSockets

4.2.4 zSeries FCP support

zSeries FCP support allows you to connect a Linux running on zSeries 800, zSeries 900, or zSeries 990 with industry standard SCSI devices. This means that you are able to configure your Linux to use a Storage Area Network (SAN) environment.

In the previous examples, a link between zLinux and IBM TotalStorage ESS is described to show how a Linux can be attached to a direct access storage device (DASD). Using a DASD can be useful if you are planning to use storage in a multipath channel environment. Multipath allows you to use multiple and simultaneous channel to access a disk.

A device configured as DASD has a capacity up to 2.347 GB. This means that if you want to use a large file system, you must configure a logical volume manager.

IBM TotalStorage ESS can be configured for DASDs, or for an open system storage devices. The open system storage configuration allows you to define the amount of bytes associated to a disk directly on IBM TotalStorage ESS systems.

The feature provided with zSeries servers is the Fiber Channel Protocol. This protocol can be configured over a FICON® card. FICON card is composed by

two different ports. Each port that is available on the zSeries server can serve only one partition to link a SCSI devices configured in an IBM TotalStorage ESS environment. Channel sharing is not implemented for this architecture. Figure 4-4 shows a typical FCP connection between a zSeries server and an IBM TotalStorage ESS.

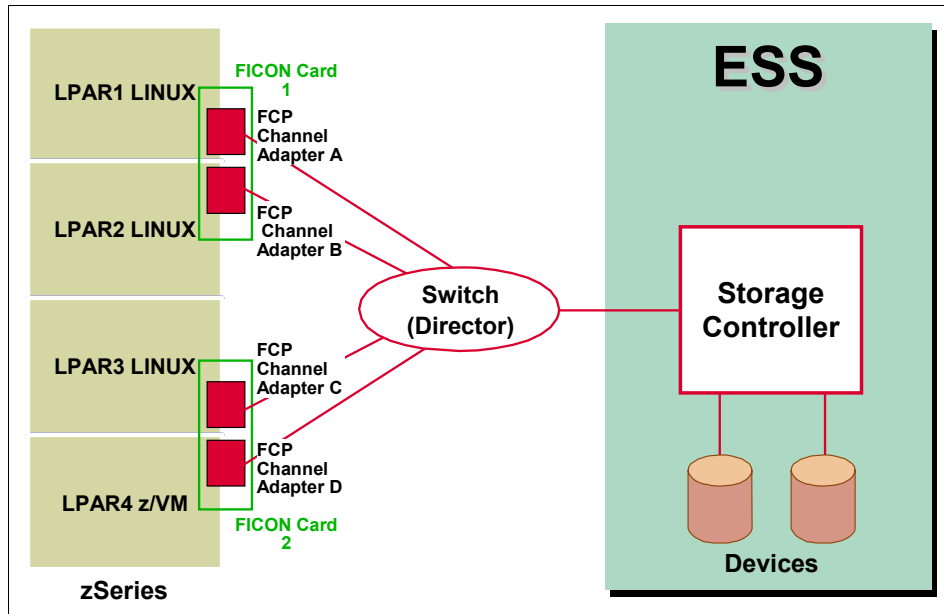


Figure 4-4 FCP connection environment

Note: You can plan to use both architectures provided with IBM TotalStorage ESS.

The feature provided with Fiber Channel Protocol allows you to connect your Linux for zSeries server farm in a SAN. This process gives you ability to use a tape library connected to the SAN, through a backup and restore server such as Tivoli Storage Manager, which is now available for Linux for zSeries.

For more information, see:

<http://www.ibm.com/servers/eserver/zseries/connectivity/fcp.html>

4.2.5 On/Off Capacity

IBM continues to enhance the On/Off Capacity on Demand (On/Off CoD) capabilities of its z990 servers by extending On/Off CoD to now include

Integrated Facility for Linux (IFL) engines in addition to the earlier support for standard CP engines that was delivered on September 15th, 2003. On/Off CoD provides a fast, non-disruptive way by which “extra” processor capacity on z990 servers can be temporarily activated when you need it – for only as long as you need it. And you pay for only the capacity you activate. For more information, see:

http://www.ibm.com/servers/eserver/zseries/announce/z990/on_off_capacity.html

4.3 All about z/VM

This section explains the value, both business and technical, of running tens to hundreds of virtual Linux systems on z/VM.

Note: *VM* means *virtual machine* not virtual memory. The virtual machine concept found in the z/VM product provides the function-rich *virtualization technology* on the zSeries server. In the context of this chapter, the virtual machine concept and virtualization technology are interchangeable. z/VM provides much more functionality than partitioning of physical resources.

zSeries Virtualization Technology as it exists today is the result of over 35 years of constant investment, innovation, and refinement. This is a proven, function-rich technology base that has benefitted the Linux on zSeries environment since day one. Customers have run business-critical solutions on VM systems since.

Virtual machine technology was introduced with the IBM System/360™, the original mainframe, introduced in 1964. At that time, VM (known as CP-67) enabled users of S/360™ systems to run multiple copies of operating systems on a single mainframe. The ability to run multiple operating systems made it possible to migrate to new releases and versions on the same mainframe system. A user could run a production workload on the old level of software, while performing the work required to migrate the environment to the new operating system.

Internally, IBM used VM to simulate the System/370™ architecture on a System/360 processor. VM would run on a real S/360 processor and create S/370™ mode virtual machines. This let IBM develop its suite of S/370-architected software before a “real” S/370 processor was available for testing.

Note: VM was also used to develop VM/XA, VM/ESA®, and the new 64-bit version of z/VM.

The ability of VM to simulate hardware and architectures today continues to give customers time-to-market advantages for their products and services.

Over the course of 30-plus years, IBM made investments in hardware, architecture, microcode, and the VM product itself to enhance the virtualization technology available with each successive line of mainframe computers. This “total system” investment is why the zSeries line of servers offers the industry's best-of-breed virtualization technology. Today, z/VM Version 4.4 is the most current VM release.

The zSeries approach to virtualization is multi-dimensional. Virtualization technology is present on the hardware, and in the z/VM software to provide a well rounded, highly efficient suite of function. An example of this would be the Start Interpretive Execution (SIE) instruction. Since any operating system can execute instructions that store the clock, change the content of registers and the like, those values must be restored before dispatching that virtual machine. To do this more efficiently, zSeries servers use the SIE instruction, which is a part of the z/Series instruction set. Since it runs directly on the hardware, it is more efficient than depending on a technique based solely in software. zSeries servers consider virtualization from a system perspective at design time. Virtualization capabilities are not added after the design as an afterthought.

On the other hand, some “virtual server” solutions consist of running multiple applications on a single operating system using virtual storage techniques to meet memory demands. z/VM, uses real resources at its disposal (processors, memory, I/O, network, etc.) to create virtual machines. Each machine can run its own independent copy of any operating system that supports the zSeries architecture (which includes Linux, of course).

Hardware architecture and facilities are presented to the “guest” system (which is “virtualized”) by z/VM in such a way that the guest operating system believes it is using a real system, not a virtual one. Figure 4-5 shows an example with multiple virtual machines, including three Linux machines: a z/OS, a VSE, and a CMS guest. The physical hardware is not seen by the guests. Instead, each guest sees the virtual machine presented by z/VM.

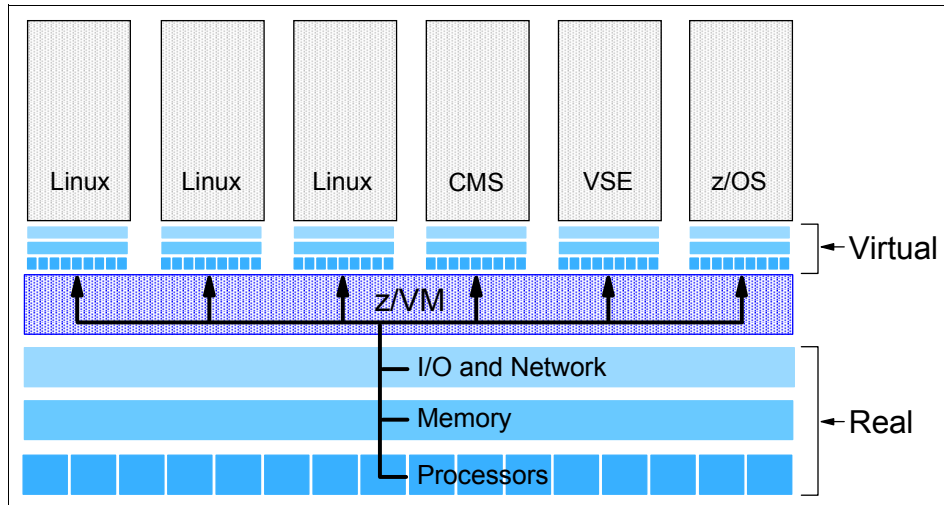


Figure 4-5 z/VM virtualization

A guest operating system can run unchanged on z/VM and not make any effort to detect whether it is running in a virtual machine. Optionally, it is possible to determine whether the “machine” is virtual. If it is, specific VM-only functions are available to guest systems.

z/VM can create a virtual machine environment that is functionally richer than a “real” environment. Examples of this include:

- ▶ Using data-in-memory techniques to transparently enhance guest system performance
- ▶ Adding virtual CPs to z/VM guests, which are seen as real CPs by the guest even when the number of virtual CPs exceeds the number of real CPs
- ▶ Simulating device and networking facilities that don't exist in the real mainframe system
- ▶ Sharing a single copy of a kernel among several guest systems

z/VM helps you to optimize resources and may provide total cost of ownership (TCO) savings, by:

- ▶ Reducing costs of the new hardware deploying virtual servers
- ▶ Fewer hardware servers can reduce costs for heating, air condition, electricity and save space requirements.
- ▶ The ability to share data and programming among multiple virtual servers can reduce maintenance complexity, effort, and expense.

- ▶ Building new virtualized servers on demand reduce time costs needed for planning a new hardware server installation.
- ▶ Providing virtual network connectivity reduce cabling and hardware costs

z/VM V4 requires G5 processor technology or better. This is found in the Multiprise 3000 server as well as the G5, G6, zSeries 800, zSeries 900, and recently announced zSeries 990. z/VM V4 (all releases) can run on the IFL and standard processor engines. IFL engines are available on the Multiprise 3000, G5, G6, zSeries 900, zSeries 990 and zSeries 800. The Linux-only model of the zSeries 800 uses IFL engines exclusively.

Note: An engine is merely another term for a processor.

4.3.1 Hardware requirements for z/VM

z/VM V4 requires at least a Multiprise 3000, G5, G6, zSeries 900, zSeries 990, or zSeries 800 processor. All z/VM V4 releases can run on an IFL as well as standard processor engines. IFL engines are available on the Multiprise 3000, G5, G6, zSeries 900, zSeries 990, and zSeries 800. The Linux-only model of the zSeries 800 uses IFL engines exclusively.

4.3.2 Virtual networking with z/VM

Virtual networks are a great way to connect guests running under z/VM. Virtual networks use z/VM virtualization to provide connectivity between guests without needing real hardware. Virtual networks enable fast and cost effective TCP/IP communication between z/VM Linux guests.

By using virtual networks, the complexity and expense of supporting the external network required to support distributed systems can be significantly reduced. The amount of cabling, hardware switches, and routers can be reduced significantly, which can reduce complexity and make networking easier to maintain and troubleshoot. Security is also improved since all communication takes place within the server and eavesdropping device (a sniffer) cannot be placed on an external cable.

Note: Connectivity in a z/VM virtual network is limited to guests running in a single z/VM image. Virtual networks cannot be used for inter-LPAR communication.

Three types of virtual networks are available to Linux guests:

- ▶ **Virtual channel-to-channel (vCTC):** vCTC networks provide point-to-point connectivity between guests without real channel allocation (as required for real channel-to-channel (CTC) connectivity). A real CTC adapter is used to connect a real mainframe to another using the channel protocol.

z/VM provides the ability to define virtual channel-to-channel adapters so users can connect virtual machines using the CTCA protocol. On z/VM, this is useful for connecting Linux virtual machines to other virtual machines that don't support IUCV, such as VSE/ESA, OS/390®, and z/OS. Virtual vCTC networks can also connect Linux virtual machines.
- ▶ **Inter-user communication vehicle (IUCV):** Point-to-point TCP/IP connections between Linux guests can be established using IUCV. IUCV is a VM-unique, virtual machine-to-virtual machine communication protocol. The Linux for S/390 and zSeries kernels include an IUCV driver, allowing you to connect two Linux virtual machines.

Linux treats IUCV connections as any other TCP/IP network connection. This results in memory-speed networking between the connected Linux servers. The speed of the IUCV “line” is a factor of processor speed. The faster your zSeries server is, the faster your IUCV network is. When you upgrade to a faster real processor, you automatically increase the speed of your virtual network. IUCV is also supported by z/VM's own TCP/IP stack, so you can use IUCV to connect Linux images to VM's TCP/IP stack.
- ▶ **VM virtual guest LAN implementation:** This type of network enables local area network (LAN) connectivity between z/VM guests. You can learn more about the details of VM guest LAN and VSWITCH in 4.3.3, “VM LAN support” on page 84.

Important: IBM recommends you to implement VM guest LAN over point-to-point for your network connectivity. It also recommends the VSWITCH architecture over router virtual machine.

Figure 4-6 shows a simple virtual network. There are three groups of five Linux servers, each connected to a virtual router using a virtual networking protocol. The virtual routers are virtual machines, each with a real networking interface to the “outside world.” In this scenario, each router is connected to a real Open Systems Adapter (OSA).

This example shows how you can connect a large number of Linux servers to a real network using fewer real network adapters, exploiting z/VM's virtual networking capabilities to essentially share the real network connections among your Linux servers. The value here is one of technology exploitation and simplicity. The virtual networking configuration looks just like something

deployed on discrete, real servers. Networking personnel do not need to learn something new when working with virtual networking on z/VM.

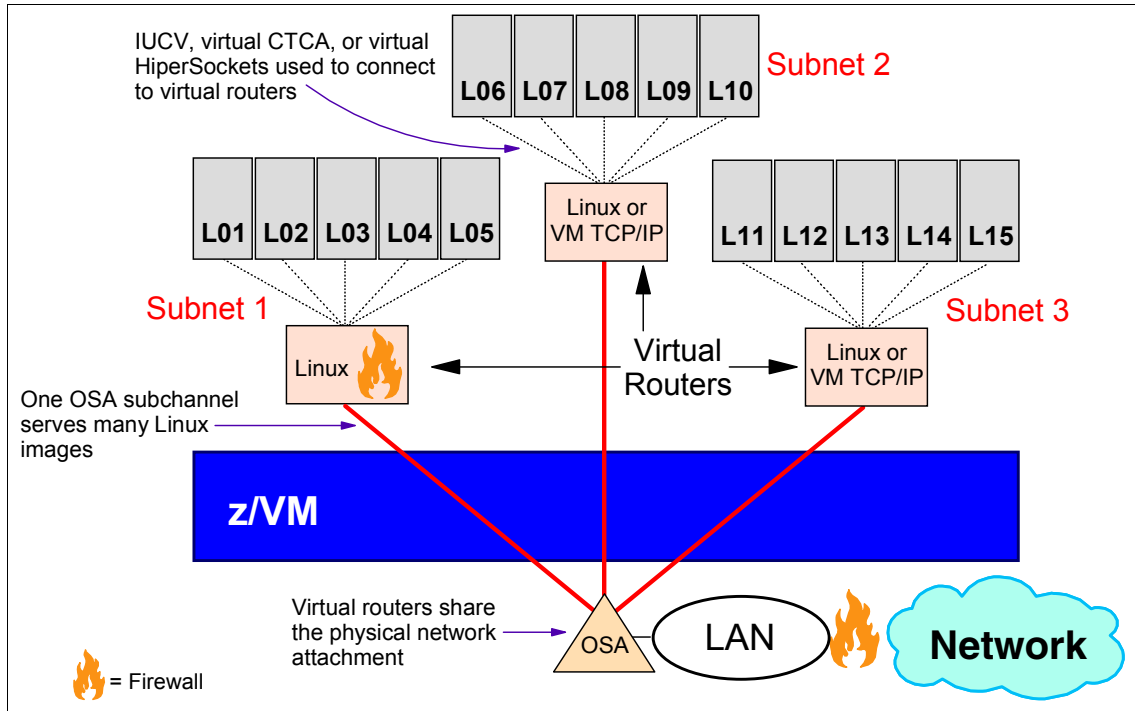


Figure 4-6 z/VM virtual networking

4.3.3 VM LAN support

This section describes the different architectures available on z/VM to implement a VM virtual guest LAN.

A VM virtual guest LAN allows the physical network to be divided into separate logical networks. These logical networks operate as if they are connected directly to a real network interface. The architectures available to implement a VM virtual guest LAN are:

- ▶ z/VM guest LAN
- ▶ z/VM virtual switch

Note: z/VM virtual switch (VSWITCH) is available only for z/VM 4.4.

z/VM guest LAN

Introduced in z/VM 4.2 and enhanced in z/VM 4.4, VM guest LAN provides VLAN networking support without relying on real LAN hardware. There are two types of VM guest LAN:

- ▶ **HiperSockets:** Introduced in z/VM 4.2, the HiperSockets VM guest LAN emulates real HiperSockets networking in a z/VM image.
- ▶ **OSA:** Open System Adapter. Introduced in z/VM 4.3, the OSA VM guest LANs emulates real OSA networking in a z/VM image.

VM guest LAN support eliminates some point-to-point LAN management challenges when using IUCV and vCTCA. z/VM is designed to enable virtual HiperSockets emulation using VM guest LANs on processors that do not support real HiperSockets such as the G5, G6, and Multiprise 3000. This enhances the virtual networking environment on those processors, while allowing customers to prepare for a real HiperSockets environment prior to moving to a zSeries server.

It is possible to define “system” guest LANs and guest LANs that are associated with a specific VM user (for example, a virtual machine). System guest LANs exist independently of any active (for example, logged-on) user. Guest LANs are associated with a user exist only while that user is active. For either type of guest LAN, authorized users can link to the LAN to participate in HiperSockets communications. There is no predefined limit on the number of virtual HiperSockets devices that you can link to a guest LAN. Nor is there a limit on the number of guest LANs that can be defined.

The virtual networking environment in Figure 4-7 shows a single Linux image (in the middle) connected to a real network interface. It is serving as a router to the “outside world” for the Linux servers connected to the three guest LANs.

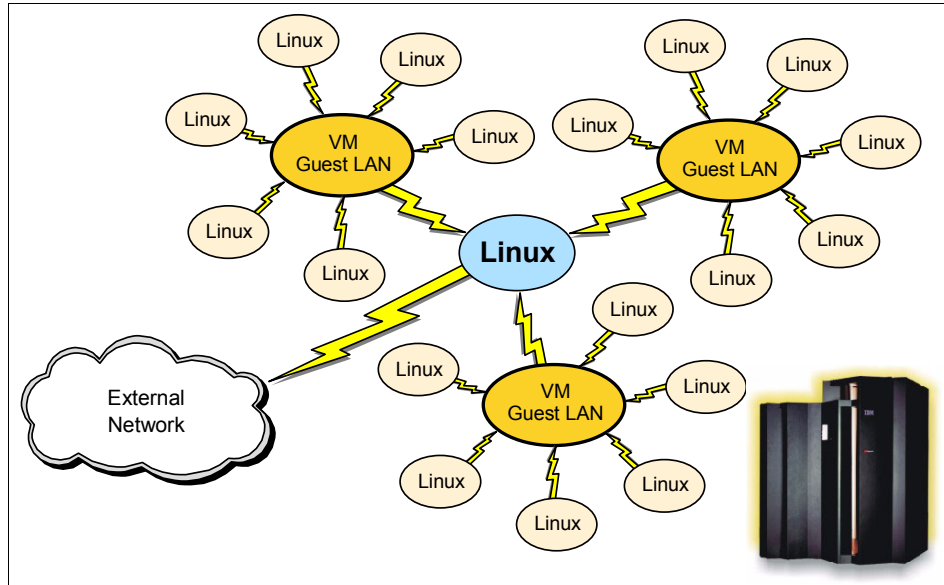


Figure 4-7 Virtual networking example

The Linux router can also transfer data from one guest LAN to another. This allows customers to isolate their Linux images on a guest LAN from Linux images on other guest LANs, using the router to secure traffic between the guest LANs.

z/VM Virtual Switch

Virtual switch technology provided with z/VM 4.4 is an extension of the guest LAN function provided in earlier releases of z/VM. Virtual Switch operates almost exactly as a z/VM QDIO guest LAN with two important exceptions:

- ▶ Direct external access link to the network through OSA Express
- ▶ IEEE 802.1Q VLAN Support

Important: IEEE 802.1Q is the VLAN standard supported under Linux and z/VM. VM guest LAN and VSWITCH supports VLAN but it is not required.

Figure 4-8 shows a network environment working in a Linux for a zSeries server farm using a VSWITCH technology. A Linux router server is not necessary to run this configuration. With VSWITCH, the TCP/IP stack functions as a controller for the OSA connection. It is not involved any data transfer activities. With the VSWITCH, the Control Program (hypervisor) directly transfer data to/from OSA to the guests, and also handles the guest to guest data transfer. The key is that the guests who are coupled to the VSWITCH are physically on the same LAN

segment as external TCP/IP stacks on the physical LAN that OSA is connected to. In fact, there is no router segmenting the LAN, and there is no virtual machine TCP/IP router overhead in the guest's communications with external servers/clients. In this case, you do not have supplementary Linux router images under your z/VM, and you can increase your performance.




Figure 4-8 VSWITCH network implementation

4.3.4 z/VM for Linux performance: Scalability

zSeries customers have many options for scaling a Linux workload on z/VM. z/VM supports vertical and horizontal growth. As mentioned earlier, z/VM exploitation of hardware technologies enhances the performance and scalability of Linux workloads.

With z/VM, you can grow your server workload vertically or horizontally. *Vertical growth* is realized when you add more system resources such as additional processing power to an existing virtual machine. You can give a Linux virtual machine more processor capacity, more virtual memory, more I/O devices and bandwidth, plus more networking bandwidth.

Horizontal growth is a typical way to grow workloads. It is easily accomplished by adding another virtual machine, or two, or three, or more on z/VM. There is an added benefit in the efficient use and sharing of system resources. This makes adding more Linux virtual servers a particularly attractive option when you consider the limited scalability of a single Linux instance.

The added scalability of Linux workloads is realized when you choose to exploit VM's data-in-memory and sharing techniques. Virtual disks in storage and minidisk cache can boost the performance of Linux servers on z/VM by making disk I/O operations execute at processor memory speeds, avoiding trips to the I/O subsystem, and waiting on spinning disks of storage. By sharing data and programming on shared, read only mini-disks server deployment time can be reduced, as well as the maintenance effort required to replicate changes to data and programming among multiple distributed servers.

z/VM support for such advanced storage functions as Peer-to-Peer Remote Copy (PPRC) and FlashCopy® benefits Linux virtual servers, although Linux itself does not support the function. Exploitation of this technology and other technologies such as FICON support is transparently passed on to the Linux servers simply because they are running on z/VM.

4.3.5 Productivity: Development and test workloads

z/VM enhances the development and testing of Linux on zSeries. z/VM makes it easier to develop, enhance, and test Linux systems. Anyone can use Linux with z/VM virtual machines. When you run Linux on z/VM, you will realize productivity gains in the areas of development and testing.

Virtual machines are easy to create, provision, modify, and decommission when you no longer need them. A z/VM environment is considerably more flexible than logical partitioning. Development and test environments are generally dynamic. They grow and shrink in size over a short period of time, which is something that is well suited for z/VM.

z/VM includes extensive built-in debug facilities. You can “instruction step” through a Linux virtual machine, interrogating storage locations, registers, and other system facilities after each instruction if you choose. You can set “traps” to stop a Linux virtual machine when a condition occurs, making it easier to target bugs.

Using z/VM virtualized system resources, you can rapidly create a complex test environment using virtual devices, networks, and facilities instead of purchasing, and waiting for real hardware to accomplish the task. With z/VM, you can conduct extreme test cases to see if your server environment can handle demand.

Virtual machines on z/VM have virtual processors. You can define anywhere from one to 64 virtual processors with z/VM. This is useful if you want to test Linux on a multiprocessor configuration, even if the underlying z/VM system is running on a single processor.

When you have your own Linux virtual machine, or two, or more, you can complete your work faster and innovation flourishes. The ability to try something on Linux in a virtual machine environment is usually a low-cost endeavor. Successful solutions are created by users who had some time and the necessary resources to experiment. Linux on z/VM gives you an environment that allows you to innovate at your best.

4.3.6 Productivity: Production workloads

Production workloads running on z/VM are easy to modify, extend, and adapt to changing business needs. Creating Linux servers on demand can help meet business needs. Time is money, especially when it comes to e-business. If business systems cannot satisfy handle the demand, or are not flexible enough to react to new opportunities, business usually suffers.

zSeries servers with z/VM can help you deal with unexpected increases in workload. You can bring Linux servers online in a matter of seconds or minutes with z/VM. You can pre-configure Linux servers offline (that is, logged off) and bring them online as soon as you need them. It is also possible to provision the server in real-time in a matter of minutes.

You can add capacity to an existing server image in real-time (without causing a server outage) using available processor capacity. You can also add real processor capacity with the zSeries CUoD and immediately add it to your existing virtual Linux server environment.

The flexibility of a virtual machine also means that Linux images can stand by as a hot backup for production servers. The attraction of this feature on z/VM is that the standby images consume minimal resources. In a discrete server world, a standby server typically costs as much as the production server it is backing up. The cost includes acquisition and support expenses as well as duplicate software licensing.

A key aspect of z/VM's value for Linux is resource sharing. This includes processor capacity, network bandwidth, and disk technology. CMS minidisks, a disk partitioning technology, can be readily shared among Linux images. This enables you to store data on one set of disks to be accessed by multiple Linux server images. A read-only minidisk can be accessed by any number of Linux virtual machines. Such a configuration is ideal for version control and upgrades of application software. With z/VM's minidisk cache support, minidisk data can

be cached using processor memory. This enables memory-to-memory data transfer rates for disk I/O to a read-only minidisk.

z/VM Resource Manager

With new release of z/VM 4.4, the Virtual Machine Resource Manager (VMRM) has been enhanced to provide the infrastructure necessary to support more extensive workload and systems resource management features that may be delivered in future releases of z/VM. To accomplish this objective, the following have been added:

- ▶ Monitor data showing actual workload achievements
- ▶ An interface to dynamically change users in workloads, workload characteristics, and goals in the VMRM configuration file without manually stopping the server
- ▶ More flexibility when managing multiple users by supporting wild-card characters on user IDs in the configuration file
- ▶ Improvements to the reliability and performance of the VMRM service virtual machine's handling of monitor data
- ▶ Serviceability enhancements including improved messages, logfile entries, and new server options

For more information about z/VM Resource Manager, and to learn more about how it can improve your Linux colony performance, see:

<http://www.vm.ibm.com/perf/>

z/VM and DASD options

Linux on zSeries can use IBM TotalStorage Enterprise Storage Server® (ESS) for disk space allocation. As shown in Figure 4-9, Linux disk space can be allocated from TDISK space, from a single disk volume, or from VM minidisks.

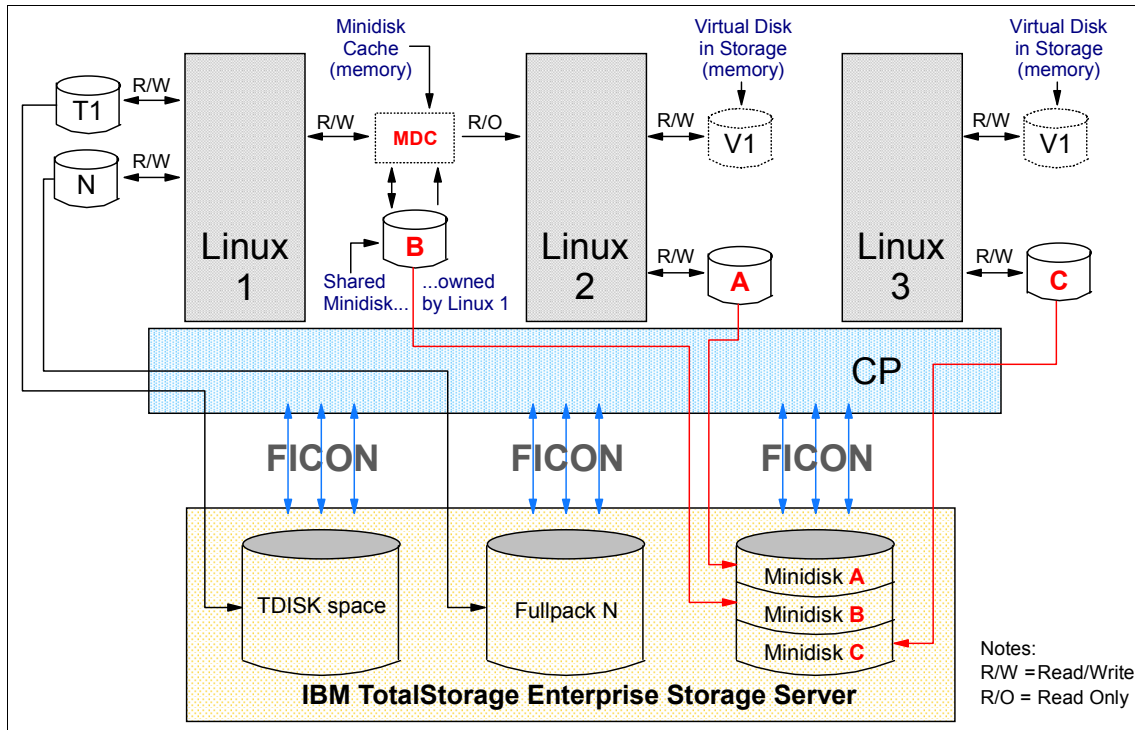


Figure 4-9 z/VM and virtual disks

4.3.7 Data backup with IBM Tivoli Storage Manager and FlashCopy

You can exploit z/VM technology for file-level Linux data backups using IBM Tivoli Storage Manager. z/VM technology improves performance of the backup operation. FlashCopy minimizes the outage time required for the backup operation.

z/VM offers a considerable amount of systems management function for Linux servers. A case in point is backup and restore using IBM Tivoli Storage Manager. Figure 4-10 represents a backup scenario that exploits z/VM's support for the FlashCopy function of the IBM TotalStorage ESS. Running Linux in a z/VM virtual machine means that you can exploit the advanced technology found in FlashCopy even though Linux itself does not support the capability.

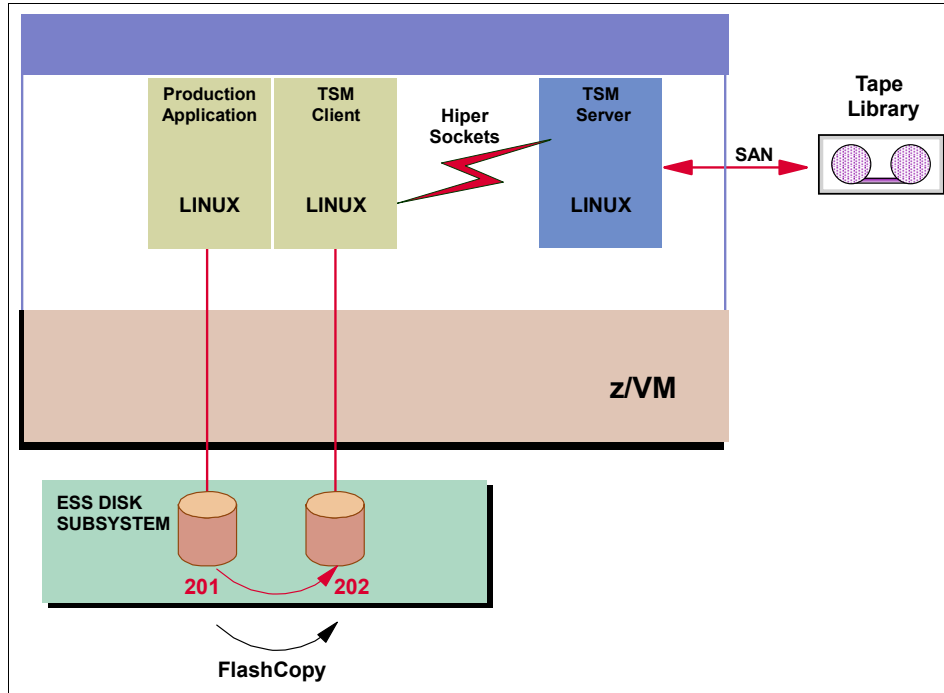


Figure 4-10 Backup and restore scenario with IBM Tivoli Storage Manager

The data that is backed up in this example is found on disk address 201. This disk is “owned” by the production application running on the left most Linux image. A FlashCopy of this disk is made using the z/VM FlashCopy command. This produces the 202 disk shown in the diagram, a duplicate copy of the 201 disk. The FlashCopy operation is a fast operation on the IBM TotalStorage ESS subsystem that minimizes the outage required of the production application running on Linux. After FlashCopy completes, the production application can resume, accessing the 201 disk as required.

The 202 disk is attached to a second Linux virtual machine that was created for the purpose of performing the backup operation. The IBM Tivoli Storage Manager client code is running in this Linux image. It sends the file-level data to the IBM Tivoli Storage Manager server that is running in a separate virtual machine, using an HiperSockets connection between the two virtual machines. The z/VM IBM Tivoli Storage Manager server, separately licensed but delivered with the z/VM product, sends the backup data to the tape subsystem through the FCP protocol providing SAN functions.

4.3.8 System management: Tivoli for Linux on zSeries

Table 4-2 shows some IBM and Tivoli solutions that you can use to manage your Linux colony under a z/Series server.

Table 4-2 System management benefits

Software	Challenges	Value
Tivoli Workload Scheduler	Multiple job scheduler	One job scheduling for z/OS and Linux servers
Tivoli Monitoring	Linux system performance problems	Centralized monitoring and control of systems
Tivoli Enterprise™ Console	Long problem resolution times and TCP/IP management	Event correlation and TCP/IP management
Tivoli Netview	Long problem resolution times and TCP/IP management	Event correlation and TCP/IP management
Tivoli Storage Manager	Long restore times to recover a Linux file	Common storage backup solution
Tivoli Configuration Manager	Long application or patch rollouts	Rollout software fixes and new applications from the central site
Tivoli Decision Support for OS/390	Unable to measure Linux and z/OS performance from a single tool	Performance and capacity analysis, capacity planning, accounting, and charge back
Tivoli Access Manager	Complex security concerns for access to a mixed z/OS, OS/390, and Linux environment	Centralized security management for OS/390 host and distributed systems
Tivoli Web Site Analyzer	Poor Web experience	Monitor health and effectiveness of Web
Tivoli Risk Manager	Web site taken down due to hacker attacks	Intrusion detection against cyber attacks
Tivoli Access Manager for e-business	Security concerns about access to key applications from the Web	Controlled Web access to key applications

Software	Challenges	Value
IBM Directory Server	Lack of comprehensive application directory	Industry-based Lightweight Directory Access Protocol (LDAP) infrastructure to provide application identify management
Tivoli Remote Control	Long problem resolution time	Remotely controlling problem systems
Tivoli NetView® for z/OS	IT staff with limited distributed tools skills	Issue Linux commands from an z/OS tool
Tivoli System Automation	Maintaining high availability	Automatically recover server outages
Tivoli Orchestrator	Unable to respond to increased demand on servers and network resources	Automatically add capacity, on demand
Tivoli Business System Manager	Unable to qualify impact of any problem	Quick root cause discovery
IBM Performance toolkit for VM	What if Linux is running on z/VM, how do you manage that environment?	Use IBM supported z/VM management tools: http://www.vm.ibm.com/perf/perfprod.html

4.3.9 Consolidation example: Server farm in a box

Many expenses (and waste) can be associated with real distributed servers. It is typical for a server image to be active only less than 20% of the time. When idle, a discrete server is unable to allocate its processor and memory resources to applications on other servers that can use the capacity.

Consolidation of distributed servers on Linux running under z/VM can offer an attractive alternative. Virtual servers on z/VM share total system resources. Processor and memory capacity is allocated to the servers that need it, when they need it.

As shown in Figure 4-11, each discrete server needs its own disk space. A lot of that space contains the same files/data found on other servers, representing a needless waste of disk space (and money). This figure shows how z/VM lets you share disk space among virtual Linux images. Not only is it an efficient use of your disk investment, but it is also a great way to provide version control for your

application software (placing your software on one disk, accessible by all your virtual Linux servers).

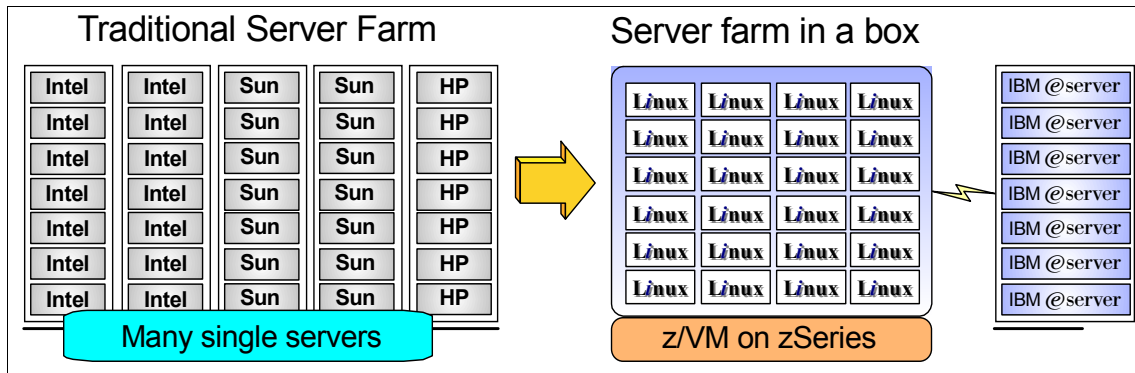


Figure 4-11 Server farm in a box

With separate servers, you need to connect real cables to each server. z/VM provides “virtual” cables for your Linux servers. These virtual cables are fast, inexpensive, very secure, and provide memory-to-memory data transfer rates.

Command and control operations require a large support staff or expensive software products. They also need corresponding dedicated servers to run the software and client software installed on each server image to communicate with the command and control servers. z/VM offers built-in command and control functions.

If required, compute-intensive server workloads can remain on discrete server images and connect to the rest of the “server farm” running on z/VM.

4.3.10 Data and application servers on a single zSeries server

Hosting your application servers on zSeries with Linux and z/VM can save you money and enhance the quality of service. Many of the applications connecting to mainframe servers are already available on Linux.

It is relatively simple to move these applications to Linux on zSeries, without requiring changes to the client community, and save money.

The server consolidation scenario in Figure 4-12 is commonly referred to as *application integration*. The cost savings are similar to the *server farm in a box* environment. The co-residency of data and application servers means that such technologies as HiperSockets can be exploited for fast connectivity between data servers and application servers without the complexities of a real network.

In Figure 4-12, the application servers running on the distributed servers (the mid-tier application servers) are replaced by Linux virtual machines running under z/VM. The real networks are replaced by the networking features offered by zSeries (HiperSockets and VM guest LANs). Linux on z/VM on zSeries facilitates the three-tier architected environment on two tiers of real hardware. There is no need to redesign the environment. The users in this scenario can continue to connect to an application server. They most likely do so without noticing that their application servers are re-hosted. That is unless they notice that their application server does not suffer the outages it used to suffer. Or, they might notice if they take a look at the cash outlay for software and middleware. Software for Linux is usually priced on a “per engine” basis. Since zSeries servers give you the ability to run large numbers of virtual Linux servers on a given engine, it is possible to reduce the charges for software substantially by reducing the number of engines.

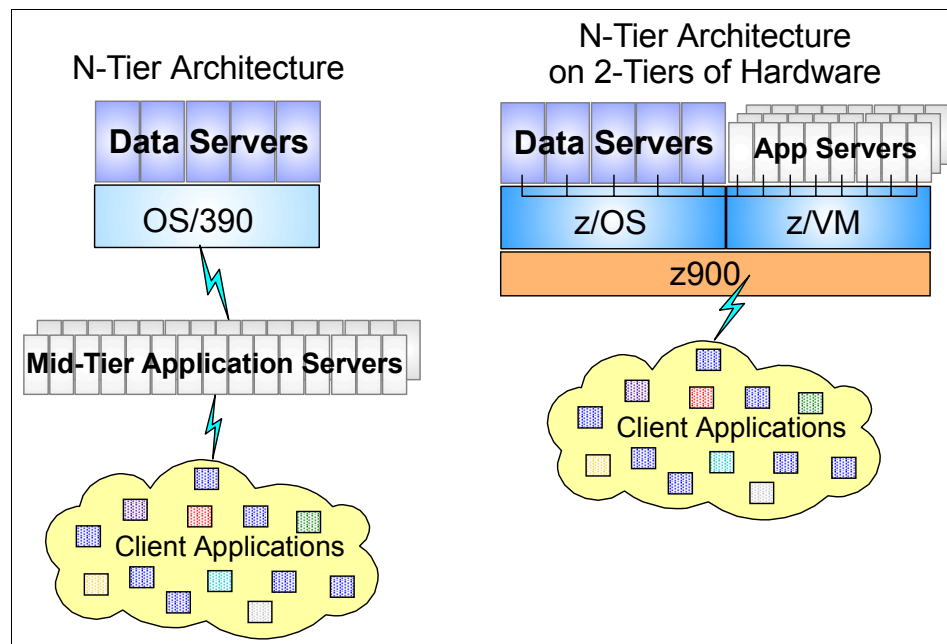


Figure 4-12 Data and application servers moving to a single zSeries server

HiperSockets enable memory-to-memory network connectivity between LPARs within the same CEC. In this diagram, an LPAR is running z/OS, and another LPAR is running z/VM with Linux application servers.

4.3.11 Horizontal growth: Adding another server

z/VM is a cost-attractive platform that handles a growing server environment. Virtualization technology is a “must have” for fast-growing server environments. You can respond to server growth with greater ease and a more efficient exploitation of hardware technologies using z/VM and Linux for zSeries. Figure 4-13 shows an example of how this works.

Sometimes growth plans are underestimated. Growth may begin with a modest number of discrete servers. Soon you find yourself adding more servers to the configuration than you anticipated. Each additional discrete server takes time to set up and bring online, contains resources that cannot be shared with other servers (memory, processors, disk space, etc.), and places demand on the people and related software products performing systems management.

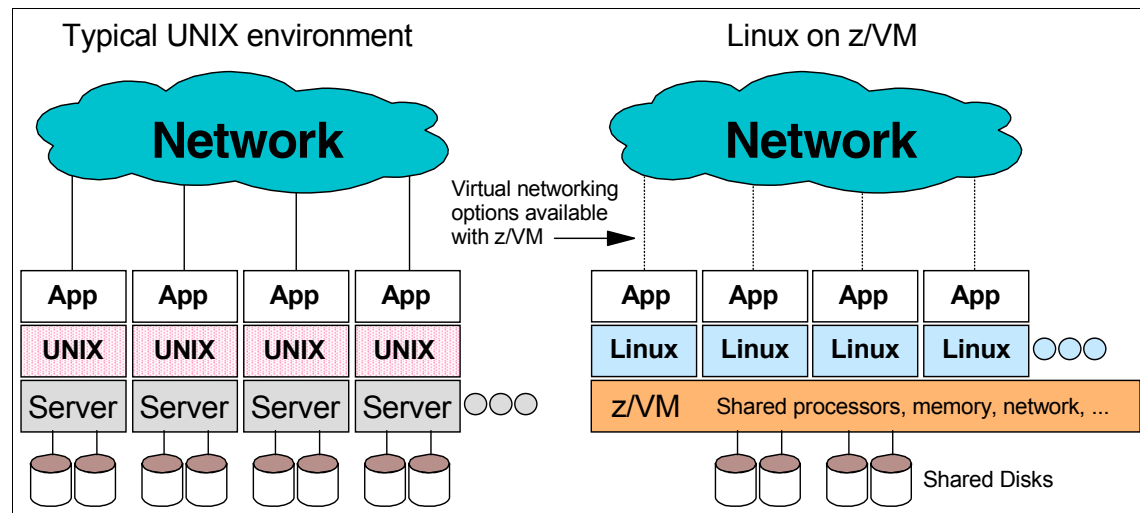


Figure 4-13 Horizontal server growth

Adding another virtual Linux server on z/VM is much less taxing on people and the IT budget. Virtual machines can be brought online in minutes, not hours or days. z/VM efficiently shares the total system resources among all the Linux server images. There is no waste of processor capacity, memory, I/O bandwidth, etc. And z/VM's built-in systems management capabilities quickly and easily adapt to the presence of another server image.

With z/VM 4.4, you can find an easy way to manage and configure virtual images. The system management application programming interface (APIs) provide a set of function that can be used from applications to allocate resources and manage servers. The APIs are intended for managing Linux virtual images, and are designed so the applications can be written by the customer.

The functions provided with APIs allow you to manage virtual images:

- ▶ Add or delete servers
- ▶ Add or remove resources
- ▶ Activate or deactivate a single image or lists
- ▶ Manage network connectivity also for the new exploitation support such as VSWITCH

Customers need to carefully consider their server growth requirements when deploying server solutions. A fast-growing server environment, although small to begin with, may be ideally suited for Linux on zSeries with z/VM.

4.3.12 Systems hardware management

This section describes z/VM 4.4 systems hardware management improvements that help to provide self-configuring, self-managing, and self-optimization facilities.

Hardware Configuration Definition

Hardware Configuration Definition (HCD) z/VM is a new facility provided with z/VM 4.4.0. It offers a comprehensive, easy-to-use I/O-configuration management environment similar to that available with the z/OS operating system. The HCD program works to help you easily create and manage the hardware and software aspects of your z/VM I/O configuration.

HCD runs in a z/VM server virtual machine. It performs the work of actually creating and changing the hardware and software aspects of your I/O configuration. It provides a user interface on your z/VM host for certain I/O configuration tasks.

HCD for z/VM can be used to create and manage the hardware aspects of the I/O configuration definition for all of the LPARs in your zSeries or S/390 server. For the software aspects, HCD for z/VM creates and manages the I/O configuration definition for only the z/VM system on which it is running. The software aspects of the I/O configuration definition for operating systems in other LPARs are handled by those operating systems.

Hardware Configuration Manager

Hardware Configuration Manager (HCM) runs on a Windows-based personal computer connected to the z/VM system through TCP/IP. HCM provides an easy-to-use graphical user interface (GUI) as well as commands to help you configure your system. You provide the needed I/O-configuration information to HCM, which processes the information and passes it to HCD.

The HCD and HCM programs work together to help you easily create and manage the hardware and software aspects of your z/VM I/O configuration. HCM provides the primary user interface to HCD for I/O configuration tasks.

The new HCM and HCD support is an easy-to-use alternative to create and manage your I/O configuration definitions compared with z/VM's existing method of using the I/O Configuration Program (IOCP) and the configuration program's dynamic-I/O configuration commands. The two methods are not intended to be used together, so you should choose one method or the other for your installation.

Figure 4-14 shows an HCM window diagram with some CTC connections.

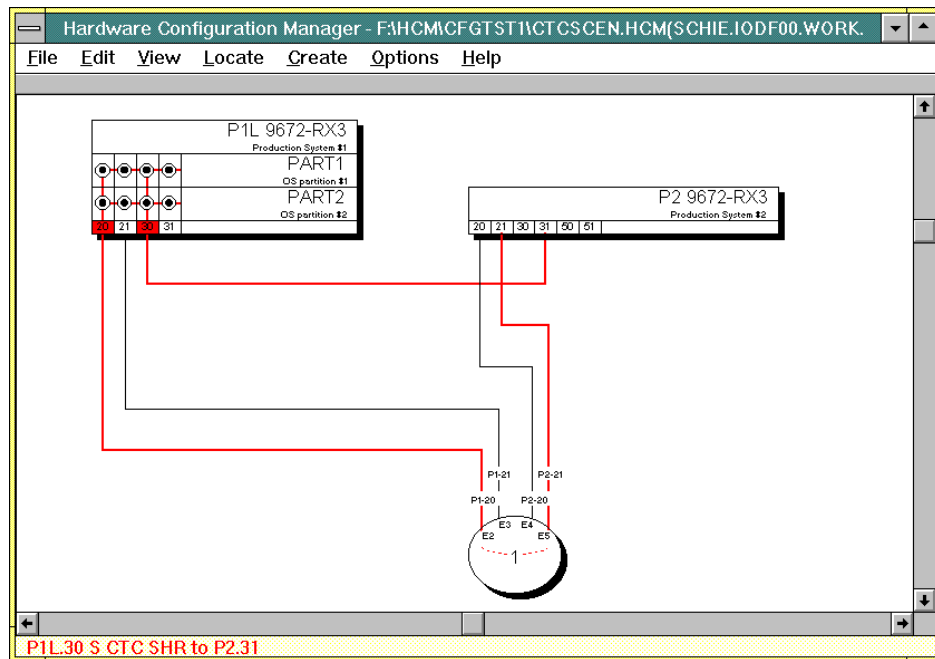


Figure 4-14 HCM diagram with CTC connections

4.4 zSeries Linux distributions

Linux distributions for zSeries are currently available from three major Linux distributors: SuSE, Turbolinux, and Red Hat.

Linux for S/390 has been generally available since early 2000 either from IBM developerWorks or from the Linux distribution partners. All the current

distributions (as of November 2002) are based on the Linux 2.4 kernel. Each distribution includes additional middleware and applications. They also include infrastructure services such as domain name server (DNS), Dynamic Host Configuration Protocol (DHCP), and Network File System (NFS) file servers, and packages such as the Apache Web server, Squid proxy server, Simple Mail Transfer Protocol (SMTP) mail server, and Samba Windows networking server. These distributions also leverage the HiperSockets technology of the zSeries server to interconnect between different partitions.

The distributions that are available for zSeries are:

- ▶ SuSE
 - SLES 7, Kernel 2.4.7
 - SLES 8, Kernel 2.4.19
- ▶ Red Hat
 - Red Hat 7.2, Kernel 2.4.9
 - Enterprise Linux AS 3, Kernel 2.4.21
- ▶ UnitedLinux 1.0 (SuSE SLES8, TurboLinux, Conectiva)
- ▶ Build your own
 - Download kernel from:
<http://www.kernel.org>
 - Download the zSeries patch and OCO modules from developerWorks

4.5 Linux application solutions

Linux distributions for zSeries contain such notable open source applications as Apache, Samba, IMAP, OpenLDAP, Perl, SendMail MTA, Squid, and PostgreSQL. Linux applications are sourced from many areas. Open source applications are available for download from the Internet. Applications can also be developed for the zSeries. Compilers are available with the Linux for zSeries distributions.

Solution providers have announced products for Linux on the zSeries server. There are more than 390 applications from more than 120 vendors available for Linux on zSeries. For a detailed product list, see:

<http://www.ibm.com/servers/eserver/zseries/solutions/s390da/linuxproduct.html>

To assist solution providers with bringing additional solutions to market, a zSeries Test Drive for Linux is available that provides remote access to Linux

installations on zSeries. The service provides access to a Linux for zSeries platform to develop, port, or test applications. For more information, refer to either of the following Web sites:

<http://www.ibm.com/servers/eserver/zseries/os/linux/lcds/>
<http://etpgw02.dfw.ibm.com/rdp.html>

Infrastructure applications

Examples of infrastructure applications announced for the zSeries Linux solutions are:

- ▶ **extend5:** This is Acucorp's integrated development environment for COBOL, which includes ACUCOBOL-GT (the flagship product). Several enterprise companies and independent software vendors (ISVs) have written or extended their COBOL code using ACUCOBOL-GT. This integrated development environment provides software developers with the ability to move their existing ACUCOBOL-GT applications to Linux for zSeries without recompiling.
- ▶ **GMx Solutions' CM_SAFE Server for zSeries:** This application provides configuration file management capabilities for open system hosts. It also provides server support for zSeries Linux installations, and agent support for all major UNIX variations and platforms.
- ▶ **zGuard:** This application from Frank Bernard Informationstechnik is a Linux-based Internet security solution for S/390 and zSeries. It has firewall, IPsec-VPN, online virus scan of various protocols including mail, HTTP, NNTP, and FTP, and basic IDS functionality.
- ▶ **Stonegate Firewall/VPN Security Software:** This application provides a unified security platform, combining firewall and VPN under a single umbrella. StoneGate is built in high-availability and reduces wide area network costs while improving reliability with Multi-Link VPN.

Line-of-business applications

Examples of line-of-business application solutions available for the zSeries Linux solutions are:

- ▶ **mySAP.com Business Suite:** This is a successor of SAP R/3 and mySAP.com solutions. It is a suite of business software and an application and integration platform designed to enable companies to manage the entire value chain across business networks, allowing the company to adapt quickly to changing business conditions.
- ▶ **LOGA 2001:** This is a modern payroll application from Personal & Informatik AG. This application incorporates remuneration in salary and wages along with travel expense accounting. A variety of solution packages can be offered,

depending on the branch involved such as the industrial sector services sector or public administration.

Personal & Informatik AG also offers a human resources management system called LOGA/ERM (Employee Relationship Management). This portal empowers employees to automate certain functions and reduces the cost for administrative work. Using the personalized home page, access is granted for applications, services, and information for daily work and career planning.

- ▶ **OpeN/2 for Linux:** This is a consolidated solution approach from S2 for authorization and payment-based transaction processing. It leverages mainframe-class computing power with reduced complexity and implementation time. OpeN/2's affinity with the Linux mainframe environment stems from its Platform Abstraction Layer (PAL®) design that isolates platform dependent code from higher-level services and applications.

Web-enablement applications

Examples of Web-enablement applications announced for the zSeries Linux solutions are:

- ▶ **eOneCommerce:** This solution from eOne Group is an e-business software application and Internet publishing tool.
- ▶ **Tamino XML Server:** This application from Software AG is a high performance information management platform based on native Extensible Markup Language (XML) storage and open standard Internet technologies. Tamino XML Server helps to find and manage any type of content across the enterprise and enables rapid implementation of robust, high-performance, and mission-critical electronic business applications based on XML standards. The server is built to:
 - Efficiently store XML documents natively, that is in their original format
 - Expose information residing in various external XML or non-XML sources (legacy data) or applications to the outside world in XML format
 - Search effectively on all information available to the Tamino application
- ▶ **DI-Atlantis:** This solution from Dimensional Insight, Inc. is a multidimensional software package designed to transform data for intuitive access and analysis. DI-Atlantis allows you to give users the information they need in a format that is easily understood.
- ▶ **GO-Global for Linux:** This product from GraphOn is a server-based solution that publishes your X applications to your network or the Web without the need for local X server software. It provides instant access to your Linux applications from any devices.
- ▶ **FrontPage Server Extension:** This application, which is from Ready-to-Run Software provides Linux support for Microsoft FrontPage Server Extensions.

The FrontPage 2002 Server Extensions SR1.1 are available for the IBM @server zSeries.

- ▶ **Coldfusion MX Enterprise:** From Macromedia MX, this application is the solution for building and deploying powerful Web applications and Web services. Using the proven tag-based scripting and built-in services in Coldfusion MX, Web application developers can easily harness the power of the Java platform without the complexity.
- ▶ **Eontec:** Financial institutions can now deploy BankFrame's EJB-based banking applications across all channels on the mainframe, meeting demands for high availability mission-critical applications and banking processes across the banking enterprise.

E-mail and collaborative software

- ▶ **IBM Lotus Domino® 6.5 Server:** This application provides a multi-platform foundation for e-mail and C&S, collaboration, and e-business. Platforms supported include Linux on zSeries, z/OS®, OS/400®, AIX®, Sun Solaris, and Microsoft Windows NT, 2000 and 2003. It supports popular client choices including Lotus Notes®, Domino Web Access with Microsoft Internet Explorer or Mozilla, Microsoft Outlook, as well as POP3 clients. Domino 6.5 offers customer flexibility and choice in hardware platform, operating system, and client access type. Industry leading security is the foundation of the Domino server. It enables administrators to manage SPAM so users do not have to. It maximizes server availability with advanced clustering, transaction logging, server fault recovery, and automated diagnostic tools.
- ▶ **Insight Server:** From Bynari, this application is an advanced e-mail server that provides a functional alternative to Microsoft Exchange. Product is built with open standards that are superior to other messaging server. Outlook clients and Web-based clients provided from Insight can access the server for all collaboration features such as contacts, calendars, and shared folders. Insight Server from Bynari facilitates complete messaging and collaboration capabilities on all IBM platforms that might provide an alternative to Microsoft's Exchange server. Support is available for a rich set of clients. The InsightConnector is an Microsoft Exchange client extension that allows MS Outlook to use Insight Server to store messages, calendar events, contacts, and other Outlook groupware functions. The Insight WebClient delivers a Web browser client that integrates with Outlook clients and provides groupware functions like calendaring, contacts, and public folders over the Web. With the Insight AddressBook client, the address book for Outlook provides functionality similar to the Exchange Global Address List.
- ▶ **CommuniGate Pro:** This application from Stalker Software offers a carrier grade messaging solution. It incorporates high performance, speed, reliability, and security with an extensive feature set. Its unique dynamic clustering architecture allows you to manage millions of accounts, while providing a

99.999% uptime requirement. It is based on open standards and supports over 24 different operating systems and hardware platforms including all IBM platforms. It provides customizable Web-based e-mail access, Internet Message Access Protocol (IMAP), SMTP/Post Office Protocol (POP), LDAP, a built-in Web server for personal Web pages, a state-of-the-art list server, multilanguage support and straightforward administration tools, and much more.

- ▶ **Samsung Contact:** This software is from Samsung SDS. Released in March 2003, it is based on Hewlett Packard (HP) OpenMail. The messaging server is an e-mail and general communication solution that runs on Linux, IBM AIX®, HP-UX, and Sun Solaris Linux. It offers the same level of access through Web, Windows, and Linux clients, and provides server level calendar and Personal Information Manager (PIM) such as handheld PC or Palm access.

Grid computing software

- ▶ **Globus Toolkit:** This product is now available for Linux on the IBM *@serverzSeries*. The Globus Toolkit for Linux on zSeries can be obtained from SuSE Linux as part of SuSE Linux Enterprise Server 8, powered by UnitedLinux or for download at <http://www.globus.org>
- ▶ **LiveCluster:** This grid middleware provided by DataSynapse is designed to solve compute and data-intensive bottlenecks and scalability constraints by harnessing the power of servers, clusters, and desktops anywhere on a network. LiveCluster allows applications running on zSeries to participate in the LiveCluster workload prioritization environment. The porting of DataSynapse's LiveCluster to support Linux for zSeries servers enables the platform to behave as an application server that delivers optimal load balancing, resource efficiency, and high availability.
- ▶ **Platform LSF:** This software provided by Platform Computing is designed to provide on-demand access to an organization's global compute resources, and balances workloads across the entire organization.
- ▶ **Platform JobScheduler:** A software solution provided by Platform Computing, which accelerates batch processing by integrating, automating, and grid-enabling silos of applications, jobs, and process flows across distributed computing clusters.
- ▶ **Platform MultiCluster:** This software provided by Platform Computing allows enterprises to create a single, cohesive computing environment with easy-to-manage resource sharing policies across geographies.

4.5.1 Mail serving scenario

Figure 4-15 illustrates a typical electronic mail application that an enterprise may use for in-house mail serving. Typically, usage of e-mail applications is sporadic, and demands on the processor are very low. Long periods of idle time are common as servers wait to be contacted by a user to send, receive, or check their mailbox. When the server is active, it usually performs I/O operations since little computation is performed for this application.

The example in Figure 4-15 makes several interesting and valid points. Regardless of the actual solution scenario, the solution-based servers are not the only servers making up the solution. Firewalls, failover protection, directory servers, backup servers, and others are typically included in production environments. Keep in mind that in most distributed solutions, any of those servers may constitute a single point of failure. Also, each server occupies floor space, consumes power, and must be monitored continuously.

The software that drives this application is uniform, and can be used by all of the servers that provide the solution. The service level required for this application is prime shift or 12 hours a day, 5 days per week.

There are two options to this scenario:

- ▶ 30 2-way Intel 1,000 MHz servers
- ▶ One zSeries 900 Integrated Facility for Linux engine

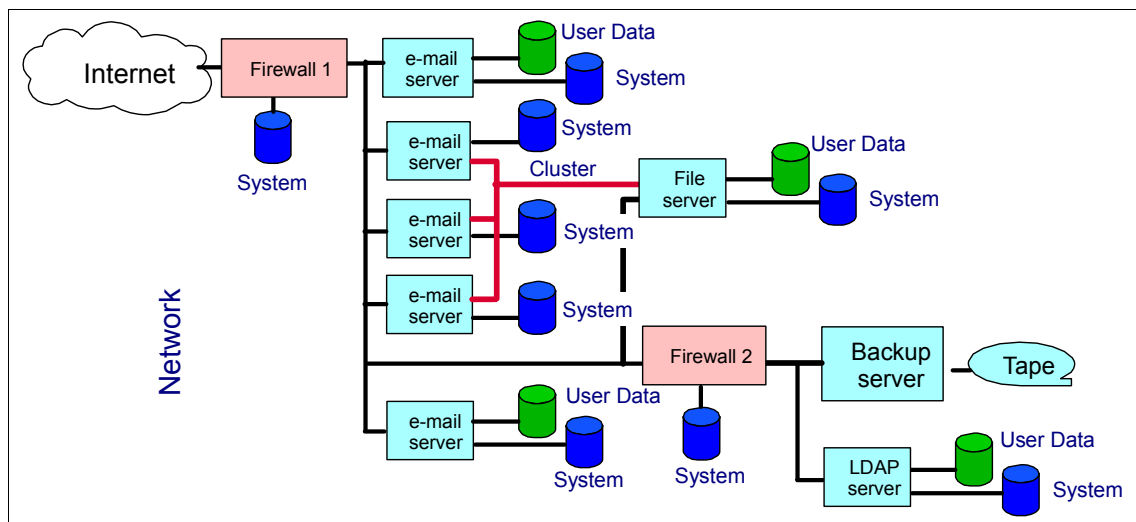


Figure 4-15 Mail serving scenario

Figure 4-15 illustrates the distributed solution. Notice that the servers are connected to a physical network which requires hardware. Also notice that each

server is running a separate copy of the operating system and application code. This is duplicated for each server in the solution, which is 30 times in this case. Each server requires disk space to support this. If additional disk storage is required for a server, the alternatives are to either purchase an additional entire disk volume and install it, or to purchase a higher capacity disk, and move the required software and data to it and install it.

Maintenance must be performed on each server individually. Some type of operator automation is required to monitor the 30 separate systems and respond to replies when necessary. Backup is done by a dedicated backup server. Client code on each system communicates with the server to perform backup operations.

Important: In the following scenario, the assumption is that the Intel servers run somewhere in the 1% to 5% utilization range.

Figure 4-16 shows the zSeries solution. Notice that the inter-server communication takes place within the machine. This allows the Linux instances to highly secure communication at memory speeds, as opposed to physical networking speeds with the usual latency, and a lesser degree of security due to exposed cabling. In addition to improving inter-server communication, much of the cost of external networking hardware and software is avoided.

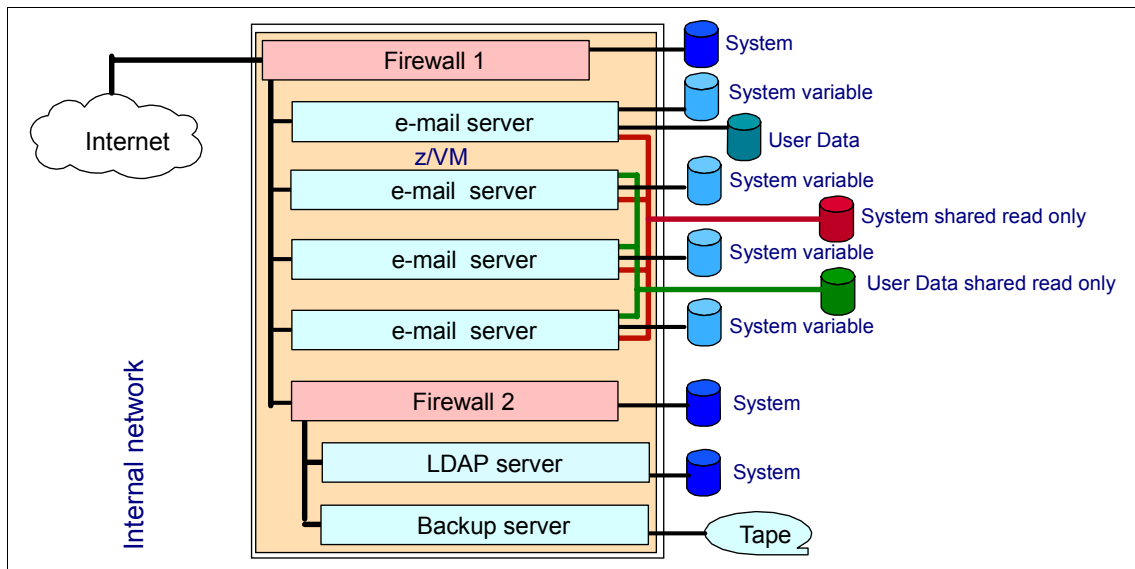


Figure 4-16 zSeries solution

Disk storage for the operating system and application code is shared by all 30 virtual servers while they must be duplicated in the distributed case. Sharing the code offers two advantages: controlling the cost of disk storage and ease of maintenance. Since only one copy of the code exists, it only has to be serviced in a single place. This results in savings in system management software as well as staffing. Although only one copy of the code is being used, 30 servers are in operation. As in the distributed case, they will issue messages, which must be monitored, and in some cases responded to. z/VM includes a programmable operator function (called PROP) that may be used for that purpose to avoid the cost of additional software to perform the task. Additional systems management software is required in complex environments.

z/VM also provides the ability to partition the disk space available to the servers. In the distributed model, increasing disk space meant buying new disk. If there is free disk space on other servers, it cannot be moved to the server needing increased disk storage. In the IBM solution, z/VM provides the ability to partition disks into minidisks that can be allocated easily to a server which requires additional disk space. Note that in e-mail applications some mailboxes grow at greater rates than others and may require expansion.

Backup is done again by a backup client. In the IBM solution, it is done by the Tivoli Systems Manager Backup server, which is pre-installed on z/VM. The Tivoli Systems Manager server communicates with clients on the virtual instances to perform backup operations.

Another factor to consider is that the IBM TotalStorage ESS disk solution is a unique and far more powerful solution than dedicated disks on stand-alone systems. The reliability characteristics of the IBM TotalStorage ESS include the advanced functions that are available in the unit such as FlashCopy. This function allows you to create a copy of a disk volume instantly. This is extremely useful for minimizing downtime for backups, or creating test data for new applications.

Cost comparison: Three-year TCO

Figure 4-17 shows an estimate of a three-year cost comparison between the distributed solution and Linux zSeries 900 solution. There are dramatic cost savings in reliability and personnel costs.




Figure 4-17 Cost comparison in the mail serving scenario

4.5.2 News group serving scenario

This case describes servers that are dedicated to news group serving such as those provided by application and Internet service providers. The application is rather simple, and demands on the processor are low. Long periods of idle time are common as servers wait to be contacted by a user that is requesting information. When the server is active, it usually performs I/O operations since little computation is performed for this application.

The software that drives this application is uniform and can be used by all servers that provide the solution. In the case of news group serving, no allowance for outages was made.

As in the previous scenario, this has three different approaches: one Intel based, one S/390 based, and one zSeries based. The alternative solutions in this case are:

- ▶ A distributed solution of 215 Intel-based servers
- ▶ 9672 six-engine Linux G6 X67
- ▶ Four zSeries 900 Integrated Facility for Linux engines

Figure 4-18 shows a logical schematic of the distributed solution. The servers are connected to a physical network that requires hardware. Also, each server is running a separate copy of the operating system and application code. This is duplicated for each server in the solution, which is 215 times in this case. Each

server requires disk space to support this. If additional disk storage is required for a server, the alternatives are to purchase an additional entire disk volume, and install it or to purchase a higher capacity disk, and move the required software and data to it and install it.

Maintenance must be performed on each server individually. Also, some type of operator automation is required to monitor the 215 separate systems and respond to replies when necessary. Backup is done by a dedicated backup server. Client code on each system communicates with the server to perform backup operations.

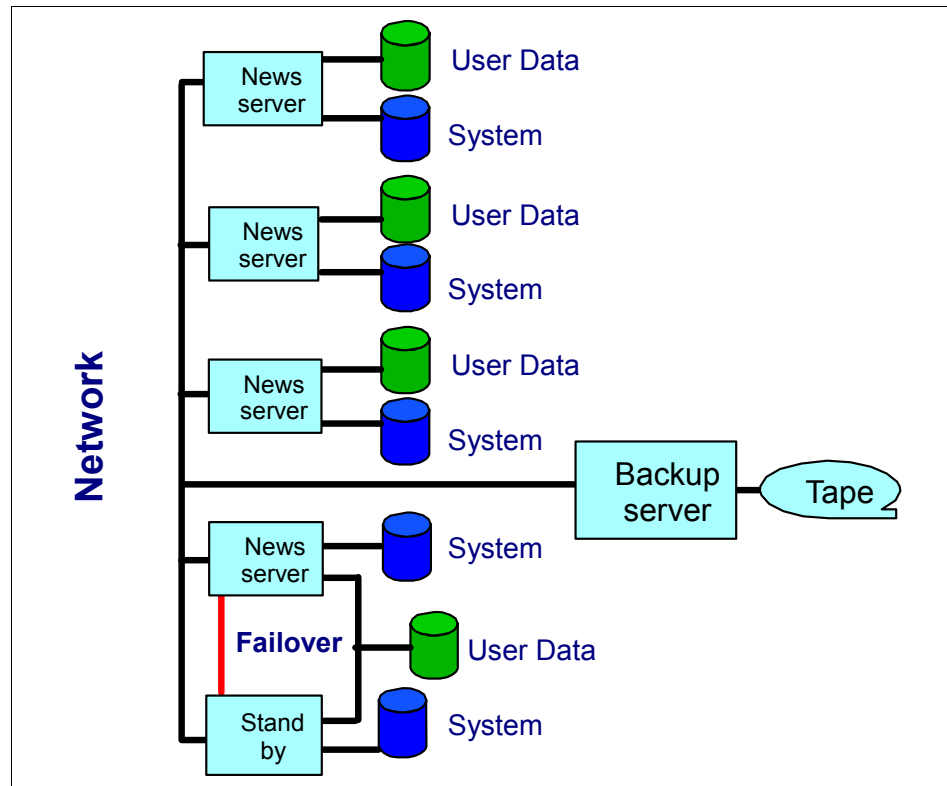


Figure 4-18 Distributed news group server scenario

Figure 4-19 shows the zSeries solution. The inter-server communication takes place within the machine, which allows the Linux instances to communicate at memory speeds, as opposed to physical networking speeds with the usual latency. In addition to improving inter-server communication, much of the cost of external networking hardware and software is avoided.

Disk storage for operating system and application code is shared by all 215 virtual servers, while they must be duplicated in the distributed case. Code sharing offers two advantages:

- ▶ Controlling the cost of disk storage
- ▶ Ease of maintenance

Since only one copy of the code exists, it only must be serviced in a single place, resulting in savings in system management software and staffing. Even though only one copy of the code is being used, 215 servers are in operation. As in the distributed case, they will issue messages that you must monitor and respond to, in some cases. z/VM's Programmable Operator (PROP) function allows you to do that, avoiding the cost of additional software to perform the task. Additional systems management software is required in complex environments.

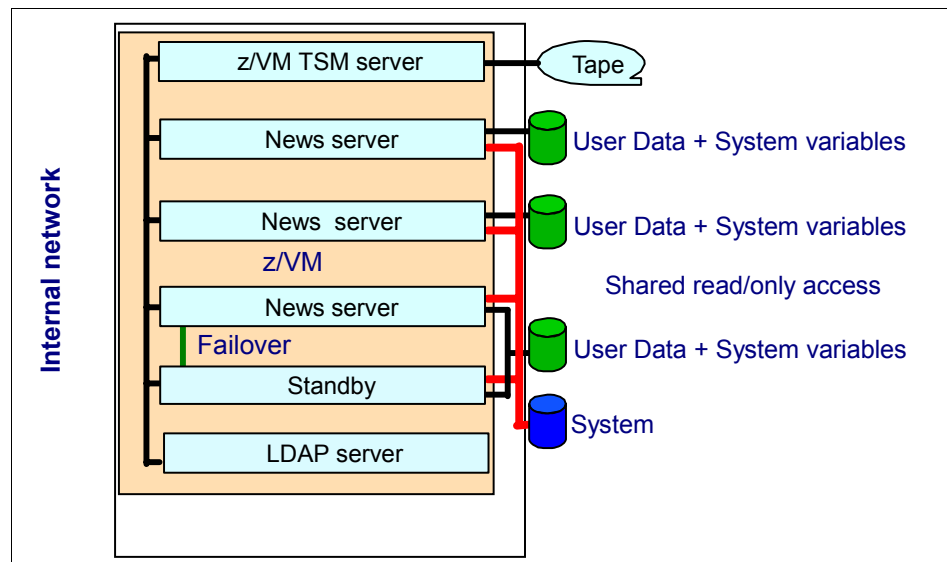


Figure 4-19 zSeries solution

z/VM can help partition the disk space available to the servers. In the distributed model, increasing disk space meant buying new disks. If free disk space is available on other servers, it cannot be moved to the server that needed increased disk storage. In the S/390 solution, z/VM can partition disks into minidisks, which can be allocated easily to a server that requires additional disk space.

Backup is done again by a backup client. In the zSeries case, it is done by the Tivoli Systems Manager Backup server, which is pre-installed on z/VM. The Tivoli Systems Manager server communicates with clients on the virtual instances to perform backup operations.

Another factor to consider is that the IBM TotalStorage ESS disk solution is unique and far more powerful than dedicated disks on stand-alone systems. The reliability characteristics of the IBM TotalStorage ESS include the advanced functions such as FlashCopy, which are available in the unit. This function allows the user to create a copy of a disk volume instantly. It is extremely useful for minimizing downtime for backups or for creating test data for new applications.

Cost comparison: Three-year TCO

Figure 4-20 shows the cost estimate for each alternative. It includes the values for the IBM solutions on zSeries and S/390. Typically, a zSeries solution requires fewer engines than a G6 solution since zSeries engines are substantially more powerful. This affects certain costs such as software and maintenance costs, which are based upon the number of engines required by a solution.




Figure 4-20 News group serving scenario comparison

4.5.3 IBM Software for Linux on zSeries

Table 4-3 through Table 4-10 on page 117 list the IBM middleware that is available to run on Linux for zSeries. Additional solutions are expected to be available in the future.

Table 4-3 zSeries and DB2 software matrix

Product name	Version	Distribution
DB2		
DB2 Connect Application Server Edition	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Connect Enterprise Edition (CEE)	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Connect Unlimited Edition (CUE)	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Intelligent Miner Modeling	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Intelligent Miner Scoring	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Intelligent Miner Visualization	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Net Search Extender	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Net.Data	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Personal Developer's Edition	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Runtime Client	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Spatial Extender	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Universal Developer's Edition	8.1	See DB2 for Linux Web site for supported platforms and distributions
DB2 Universal Database Enterprise Server Edition	8.1	See DB2 for Linux Web site for supported platforms and distributions

Table 4-4 zSeries and IBM Software

Product name	Version Distribution	
IBM		
IBM Application Workload Modeler for Linux on zSeries	1.1	Red Hat Linux 7.2 - SLES 7
IBM Application Workload Modeler for Linux on zSeries	R1	Red Hat Linux 7.2 - SLES 7
IBM CICS Transaction Gateway	5	SLES 7 - Turbolinux V6 - Turbolinux V6.5
IBM Developer Kit for Linux, Java 2 Technology Edition Java Technology development environment	1.4	See the developerWorks Web site for the supported platforms and distributions
IBM Developer Kit for Linux, Java 2 Technology Edition Java Technology development environment	1.3.1	See the developerWorks Web site for the supported platforms and distributions
IBM Directory Integrator	5.1	Red Hat 7.1, or later
IBM Directory Server	5.1	United Linux 1.0 - Red Hat 7.1, 7.2 SuSE 7.0 TurboLinux 6.5
IBM Directory Server	v4.1	Redhat 7.2 - Turbolinux 6.5 - SuSE Linux 7.0
IBM Directory Integrator	5.1	Red Hat 7.2 SuSE with kernel level of 2.4
IBM eServer Integrated Platform for e-business on zSeries 1.1	1.1	SLES 7 provided
IBM HTTP Server	1.3.19.3	
IBM Object REXX	2.2	SLES 7
IBM Screen Customizer	2.0.7	
IBM Screen Customizer	2.0.6	Client on a Host-on-Demand Server: Linux on zSeries
IBM SecureWay Directory	3.2.2	RedHat 7.2 SuSE 7.2
IBM VisualAge Generator Server for Linux	4.5	SUSE 2.4.7 kernel - Enterprise Edition 7
IBM VSE/ESA	2.7	

Table 4-5 zSeries and IMS™ software matrix

Product name		Version Distribution
IMS Note: IMS-Connect only supports Linux		
Information Management System (IMS) Connect	2.1	Note: IMS Connect running under z/OS or OS/390 can be used by an application running on Linux S/ 390 or zSeries to access local applications and data residing on IMS
Information Management System (IMS) Connect	1.1 1.2	Note: IMS Connect running under z/OS or OS/390 can be used by an application running on Linux s/ 390 or zSeries to access local applications and data residing on IMS

Table 4-6 zSeries and Informix® software matrix

Product name	Version Distribution	
Informix		
Informix 4GL	7.32	Minimum: kernel 2.4.2-50 gcc 2.95.3 glibc 2.2.2-10
Informix C-ISAM	7.24	Kernel 2.2.16 (GLIBC 2.1.3) - Linux for s/390
Informix Client SDK	2.81	
Informix Connect Runtime	2.81	
Informix Dynamic Server Extended Edition and Extended Edition Unlimited	9.4	Minimum of 2.4.x kernel
Informix Server Administrator	1.5	
Informix SQL	7.32	Minimum: kernel 2.4.2-50 gcc 2.95.3 glibc 2.2.2-10

Table 4-7 zSeries and U2 software matrix

Product name	Version Distribution	
U2		
IBM UniData	5.2x	SLES 7

Table 4-8 zSeries IBM Tivoli software matrix

Product name	Version	Distribution
Tivoli		
Tivoli Access Manager for e-business	5.1	SLES 8
Tivoli Access Manager for e-business	4.1	SLES 7 31bit - SLES 7 64 bit
Tivoli Access Manager for e-business	3.9	SLES 7
Tivoli Access Manager for OS	5.1	SLES 8
Tivoli Access Manager for OS	4.1	SLES 7 - Red Hat 7.2
Tivoli Configuration Manager	4.2	SLES 7 31bit - SLES 7 64bit
Tivoli Decision Support for OS/390 - System Performance	1.6.0	SLES 7
Tivoli Decision Support for OS/390	1.5.1	SLES 7 - TurboLinux 6.5
Tivoli Directory Integrator	5.1	Red Hat or SuSE with kernel 2.4.x
Tivoli Directory Server	5.2	Red Hat ES 3.0 - SLES 8
Tivoli Distributed Monitoring	4.1	SLES 7
Tivoli Enterprise Console	3.9	Red Hat 7.2 - SLES 7 - SLES 8
Tivoli Enterprise Console	3.8	SLES 7
Tivoli Enterprise Console	3.7.1	SLES 7
Tivoli Identity Manager	1.1	SLES 7 (Client Only)
Tivoli Intelligent ThinkDynamic Orchestrator	1.1	SLES 7
Tivoli Management Framework	4.1	Red Hat 7.2 - SLES 7
Tivoli Monitoring	5.1.1	SLES 7 - Red Hat Server 6
Tivoli Monitoring	5.1	SLES 7
IBM Tivoli Monitoring for Business Integration	5.1.1	Red Hat Enterprise Linux 2.1 - SLES 7
Tivoli Monitoring for Transaction Performance	5.1	Red Hat Linux 7.1 - TurboLinux Server 7 - SLES 7
IBM Tivoli Monitoring for Web Infrastructure	5.1.2	Red Hat 7.2 - SLES 7 (WAS5 only)
Tivoli Netview for z/OS MSM Agents (Client Only)	5.1	Red Hat 7.2 - SLES 7
NMC Server (Server Only)		
Tivoli Provisioning Manager	1.1	SLES 7
Tivoli Privacy Manager for e-business	1.2	SLES 7
Tivoli Remote Control	3.8	SLES 7
Tivoli Risk Manager	4.2	SLES 8
Tivoli Risk Manager	4.1	SLES 7 (Client Only)
Tivoli Risk Manager	3.8	SLES 7 (Client Only)
Tivoli Software Distribution	4	SLES 7
Tivoli Storage Manager	5.2	SLES 7, SLES 8
Tivoli Storage Manager	5.1.5	SLES 7 (Client Only)
Tivoli Storage Manager	5.1	SLES 7 (Client Only)
Tivoli Storage Manager Client	4.2	SLES 7
Tivoli Switch Analyzer	1.2	
Tivoli System Automation for Linux	1.1	SLES 7, SLES 8
Tivoli User Admin	3.8	SLES 7 (Client Only)
Tivoli Workload Scheduler	8.2	Red Hat 7.2 - SLES 7

Table 4-9 zSeries IBM WebSphere software matrix

Product name	Version	Distribution
WebSphere		
WebSphere Application Server	5.1	UnitedLinux 1.0
WebSphere Application Server	5.0.2	Red Hat 7.2 - SLES 7 - UnitedLinux 1.0
WebSphere Application Server	5.0.1	Red Hat 7.2 - SLES 7
WebSphere Application Server	5	SLES 7 - Red Hat 7.2
WebSphere Application Server Advanced Developer Edition for Linux	4	SuSE 7.0 - Turbolinux Server 6.5
WebSphere Application Server Advanced Edition for Linux	4	SuSE 7.0 - Turbolinux Server 6.5
WebSphere Application Server Advanced Single Server Edition for Linux	4	SLES 7
WebSphere Application Server Advanced Edition for Linux	3.5	SLES 7 - TurboLinux Server 6.0
WebSphere Application Server for Developers	5	SLES 7 - Red Hat 7.2
WebSphere Application Server Enterprise	5.0.1	SLES 7
WebSphere Application Server Enterprise	5	SLES with 2.4 Kernel
WebSphere Application Server Enterprise for Developers	5	SLES with 2.4 Kernel
WebSphere Application Server Network Deployment	5.1	UnitedLinux 1.0
WebSphere Application Server Network Deployment	5.0.2	Red Hat 7.2 SLES 7 UnitedLinux 1.0
WebSphere Application Server Network Deployment	5	SLES 7 - Red Hat 7.2 (Application Server Only)
WebSphere Commerce Business Edition for Linux	5.4	SLES 7
WebSphere Business Integration - Event Broker	5	SLES 8
WebSphere Commerce Business Edition for Linux	5.5	Red Hat 7.2
WebSphere Commerce Professional Edition for Linux	5.5	Red Hat 7.2
WebSphere Host Access Transformation Services	5	Server inherits WebSphere Application Server requirements; client requires supported Web browser
WebSphere Host on Demand	8	
WebSphere Host on Demand	7	
WebSphere Host on Demand	6	Server: Linux on zSeries

Table 4-10 zSeries IBM WebSphere software matrix (continued)

Product name	Version	Distribution
WebSphere		
WebSphere MQ	5.3	Kernel 2.4, GLIBC 2.2.4 or higher
WebSphere MQ Everyplace	2	Linux running kernel 2.4 for zSeries
WebSphere MQ Everyplace	1.2	Linux Distributions running Kernel 2.4
WebSphere MQ for Linux	5.3	2.4 Kernel with glibc 2.1 or later
WebSphere Personalization Server for Multiplatforms	4	Red Hat 7.1 or 7.2 - SLES 7
WebSphere Personalization Server	3.5	
WebSphere Portal for Multiplatforms Enable	5	SLES 7
WebSphere Portal for Multiplatforms Extend	5	SLES 7 (not support all components)
WebSphere Portal Server for Multiplatforms	4.2	SLES 7
WebSphere Portal Server for Multiplatforms	4.1	SLES 7
WebSphere Studio Application Monitor for z/OS	2.1	SLES 7 (Managing Server)

Table 4-11 zSeries Lotus software matrix

Product name	Version	Distribution
Lotus		
Lotus Domino	6.5	UnitedLinux 1.0 SP2
Lotus Domino Web Access	6.5	UnitedLinux 1.0 SP2 (Server Only)
Lotus Workplace Web Content Management	1.1	

Table 4-12 zSeries Rational® Software matrix

Product name	Version	Distribution
Rational		
Rational ClearCase	5	SLES 7
Rational ClearCase CT	5	SLES 7
Rational ClearCase MultiSite	5	SLES 7

4.6 How IBM customers are using Linux

The zSeries server is widely recognized as an integrated business server that scales both vertically and horizontally. It is reliable, scalable, and recognized as one of the most flexible, easy to use systems in the industry. It can run multiple environments and help to quickly deploy applications.

These attributes position zSeries as one of the best platforms to manage the complexity and cost of e-business enablement. Key characteristics of Linux on zSeries, such as a new generations of applications, integration, and consolidation, strongly support the IBM initiatives. They can result in measurable customer benefits for the deployment of e-business solutions.

4.6.1 Customer scenario with SAP on zSeries

A European company decided to improve the availability and cost-efficiency of the technology platform for its distributed SAP R/3 enterprise resource planning applications by consolidating them onto Linux. The company upgraded its SAP implementation onto an IBM @server zSeries 990 mainframe with multiple logical partitions (LPARs) running IBM z/OS (for the database server), IBM z/VM and SuSE Linux for the application servers.

The zSeries server, which leverages the Hipersocket feature for simpler use and implementation, runs the following logical partitions (LPARs):

- ▶ Six LPARs running z/OS V1.3 act as the SAP database servers (all SAP data is stored in an IBM DB2 database).
- ▶ Another six LPARs run z/VM V4 and SuSE Linux Enterprise Server (SLES) V8, powered by UnitedLinux, and serve as the SAP application platform.
- ▶ The test LPARs also runs SuSE Linux Enterprise Server (SLES) V8, powered by UnitedLinux.

The configuration is shown in Figure 4-21.

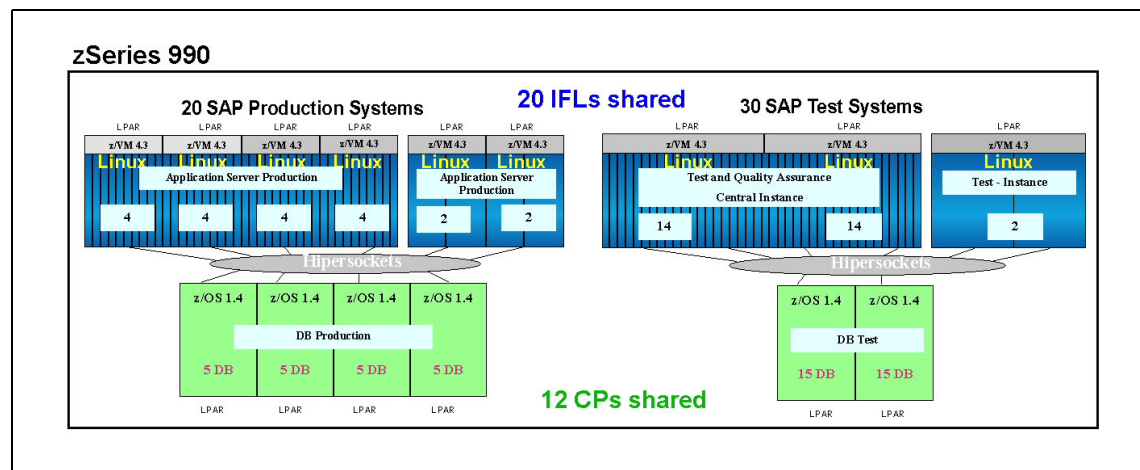


Figure 4-21 SAP configuration scenario

IBM Global Services - Integration Technology Services (ITS) provided the hardware and Linux implementation services as well as z/VM, z/OS, and LPAR configuration. ITS will continue to provide z/VM, z/OS, and Linux support for the SAP platform with IBM specialists onsite. The SAP system currently supports 1,000 users, and is configured to support an expected 1,500 users.

The customer decided to implement SAP on zSeries using Linux for the following reasons:

- ▶ High availability and higher scalability of the systems
- ▶ Higher performance of batch processing
- ▶ Better disaster recover
- ▶ Fast access to DB2 through Hipersockets
- ▶ No external network (-management) required
 - Reduce complexity/points of failures and maintenance effort
- ▶ Creation of new SAP systems is quick and very easy (z/VM cloning)
 - Server provisioning on demand
- ▶ Cost of ownership (TCO) advantages through
 - Only one platform
 - Reduced administration effort
 - Reduced operational costs
 - Reduced staff
 - Simplified backup concept

The company has realized sizable savings of time and resources because the IBM/Linux solution is easier to use and manage, and has greater uptime. The company has also saved money by using Linux instead of proprietary solutions.

4.6.2 Local Government, Italy

A local government agency in the center of Italy that provides services to citizens has implemented zSeries Linux solutions. The central Italian region of Lazio consists of the cities of Rome, Frosinone, Latina, Rieti, and Viterbo.

For five years, the regional web site has been driven by Lotus Domino on two Microsoft Windows servers and provided a wealth of information to the citizens, businesses, and municipalities. The regional government decided to transform its Web site into an e-government portal to deliver information and provide self-service applications for the people and businesses. It also connected the regional government to all its municipalities.

Approximately two years ago, the agency established a corporation with it's, IT arm to spearhead the region's e-government effort. With a staff of 20 seasoned IT professionals, the corporation relied on the e-business expertise of IBM Global Services to ensure delivery of some of the important e-government applications including the portal:

- ▶ Integrated Technology Services provided capacity planning and infrastructure design. It delivered the solution to build an integrated Linux for zSeries platform.
- ▶ Business Integration Services provided the design and architected the new portal.

The solution is an e-government Web site running under Linux that integrates the content and back-end applications and enables inter-organizational workflow. The portal also provides self-service applications for the community, supporting its sanitary procedures and services. The inter-organizational application provides services between government agencies.

The key components used for the solution are:

- ▶ Software
 - z/VM 4.3
 - IBM WebSphere Application Server, Version 5.0.2
 - IBM WebSphere Network Deployment, Version 5.0.2
 - IBM DB2 Universal Database Version 8
 - IBM Lotus Domino
 - IBM HTTP Server
 - Oracle Database Server 9i
 - SuSE SLES 8, 31 bit
 - SuSE SLES 8, 64 bit
- ▶ Servers
 - Existing IBM mainframe 9672
 - New IBM zSeries 800
 - IBM TotalStorage ESS
- ▶ Services
 - IBM Global Services
 - Integrated Technology Services
 - Business Integration Services

The infrastructure was proposed in two ways:

- ▶ A traditional LPAR running infrastructure components
- ▶ z/VM to virtualize front-end applications with the IBM HTTP server

To complete the infrastructure, the agency decided to use an IBM TotalStorage ESS to provide storage support. The IBM TotalStorage ESS system is connected in a different way to provide native and integrated capacity to Linux and z/VM operating system. It also provides large data space allocation for the database server systems through a connection to a SAN.

For the complete story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5DMV8T?OpenDocument&Site=eserverzseries>

4.6.3 The largest z990 for a Linux solution

Recently, an IBM customer purchased the largest z990 ever sold in history for a Linux z/VM workload. It is a Model 316 with one CP dedicated to their z/OS workload, and 15 IFLs dedicated to their Linux Collaboration workload. Also clothed with the z990 is an ESS Model 800 upgrade, a 3494 ATL upgrade, and a Brocade M12 SAN Switch.

IBM anticipates the z990 Linux z/VM system to grow to as much as 30 to 60 IFLs over the next two years to run the e-mail, voicemail, calendaring, and Blackberry Oracle Collaboration Suite workload for 260,000 users across all of customer departments, as well as a major ERP system rollout.

The irony here is that the customer's strategy for this project originally was to move off, once and for all, their old PROFS/VM system to something more modern with two SUN 85-way systems. So where do they go? To a z990 mainframe with z/VM running with Linux.

Another important point is this sale was the result of successful collaboration of an IBM Business Partner VICOM/Infinity Systems, IBM Inside Sales, IBM Linux Impact Team, IBM zSeries Linux Sales, IBM ATSS, IBM Pricers, IBM Lab's, and IBM FTSS.

For the complete story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5KLQ57?OpenDocument&Site=eserverzseries>

4.7 zSeries services solutions

IBM Global Services offers many different services for zSeries and Linux solutions. The following sections provide a glance at the available offerings. For more details, see the IBM Global Services Web site at:

http://www.ibm.com/services/e-business/linux_2.html

Operational support services

Support Line for Linux provides a comprehensive around-the-clock enterprise-level remote usage and defect support for major distributions of the

Linux operating system. It also provides such support for all IBM and many non-IBM applications that operate in a Linux environment.

IBM Global Services portfolio of services

IBM Global Services offers a comprehensive portfolio of Linux services including:

- ▶ **IBM Migration Services: Consolidates file/print and Web serving workloads to Linux for zSeries:** This service helps customers to consolidate workloads from distributed server farms by providing an assessment of which workloads can be consolidated onto Linux on zSeries. Then it completes the consolidation for them.
- ▶ **IBM Migration Services: Sendmail Advanced Message Server for Linux:** This service migrates a customer's existing sendmail servers and users (on Windows, UNIX, or Linux servers) to Sendmail Advanced Message Server (SAMS) running on Linux for zSeries.
- ▶ **IBM Migration Services: Bynari Insight Server for Linux:** This service migrates a customer's existing Windows Exchange servers and users to Bynari Insight Server running on Linux for zSeries.

IBM Global Services also offers locally delivered services including:

- ▶ Linux Rapid Deployment Solution for installing Linux; installing, setting up, and enabling TCP/IP; and installing and configuring Samba or Apache. This service enables a customer to expedite the deployment of applications on their Linux for zSeries server.
- ▶ Middleware, database, and connector installation and configuration services for WebSphere Application Server Advanced Edition, DB2, MQSeries, etc.

Linux Strategy Workshop

This offering includes two days of information gathering at the customer site, a one-day workshop activity, and one day to finalize the conclusions. Based on the customer interviews during the first two days, a tailored Linux strategy is developed with the decision makers and IT management personnel. At the end, the customer receives a high-level, but customized Linux strategy. The decision makers and IT management personnel also receive an overview of what Linux can do for them, and the cost efficiencies gained by using Linux.

Linux Solution for e-business

The Linux Solution for e-business service provides rapid deployment of a customer's Linux and zSeries-based e-business environment. This solution is tailored to fit the specific needs of the customers by allowing them to select a wide variety of middleware and applications, which will be installed and configured on a base of an IBM-supported Linux operating system.

VM and Linux installation support

Beginning with the LPAR configuration of the zSeries, this offering helps customers to implement Linux in virtual machines running under VM. The service includes installation of VM and Linux as a VM guest.

Linux high availability for zSeries installation support

This specialized offering implements a high availability solution using Linux on zSeries.

Bynari installation support

Bynari Insight Server for Linux services provides a compelling cost and systems management proposition for customers to consolidate messaging workload from their Windows NT server farms to multiple Linux images running on a single IBM @server.

Customers often have Microsoft Exchange as their messaging server platform. IBM Migration Services can provide services to migrate to Bynari Insight Server for Linux. This results in server consolidation of those workloads that can minimize the TCO, enable system management capabilities, and enhance application availability.

Migration to Sendmail Advanced Message Server for Linux

The primary target of this offering is customers who are already using sendmail on UNIX-like platforms or on multiple servers. During the project, all mail servers are consolidated onto one central Sendmail Advanced Message Server running on Linux on zSeries.

DB2 Connect installation support

DB2 Connect enables the connection of applications to the mainframe databases. This service sets up DB2 Connect to make the host data directly available to clients as PCs or network clients.

Workload and server consolidation

Workload and server consolidation is one of the major driving factors in the IT industry. The zSeries is a target platform for consolidation issues. Therefore, IBM Global Services offers several solutions for customers that need to consolidate. Some of the available services are:

- ▶ **Consolidation for file/print services and Web serving workloads**
- ▶ **Migration services for DB2 Universal Database for S/390:** This offering includes a migration to DB2 Version 7 on OS/390 and DB2 for Linux running under a system image (LPAR), a guest system (VM), or stand-alone Linux.

The project includes installation, customization, and verification, plus basic skill training for operations staff.

- ▶ **SmoothStart™ Services for MQSeries:** These standardized projects contain installation planning and on-site project management to ensure a quick and easy installation of MQSeries. These projects also offer a basic skill training for staff.

IBM Managed Hosting: Linux Virtual Services

This service delivers capacity as an on demand utility service, giving the company a stable infrastructure for e-business on demand, and allowing it to pay only for the processing, storage, and networking capacity it needs at any given moment. The service allows you to increase your enterprise reliability, flexibility, and scalability without up-front capital investment.

4.8 Why Linux for zSeries is on demand

Optimizing resources and operations is the goal for a company that is going to dynamically respond to needs of customers, employees, or partners. An on demand business is required to respond with speed to any customer requests or market opportunity.

To allow the new market requests, a different conceptual operating environment is needed. To have an on demand operating environment, your system must be:

- ▶ Integrated
- ▶ Built with open standards
- ▶ Virtualized
- ▶ Autonomic

These requirements are fundamental to be prepared for changes and to be capable to adopt flexibility in your business processes. A solution provided with Linux for zSeries under z/VM can achieve this requirement and add more value to your business. Many of these concepts are already described in this chapter, but do not focus on your business view.

A typical environment with Linux for zSeries is built under z/VM to *virtualize* images that will be used for a Web application. Horizontal scaling of your server means to *virtualize* images under z/VM to serve specific functions. HTTP server is necessary to serve Web pages. Usually such pages are built over an application server such as WebSphere. These applications use an LDAP server to manage user security, and a database server to read or write data for each transaction. Each service works in separate virtualized images. If your systems need more resources, you can add them using Capacity on Demand provided by zSeries, and add new images for a horizontal growing with z/VM.

Linux for zSeries servers are designed to provide a highly secure yet *integrated* environment in which applications and data can share system resources. zSeries servers give you the choice to run multiple Linux, z/OS, z/VM, TPF, or VSE/ESA systems simultaneously, while *integrating* key middleware components. The level of flexibility means that you can *integrate* your business with the zSeries while helping to maintain an effective mix of platforms and applications to achieve your goals.

If you have used IBM software for your environment, you have built an *open standard* environment. Linux already uses open technologies and IBM products can be associated to them. For example, WebSphere Application Server is based on Java technology, which is becoming the new standard programming language. An application written with the Java standard can be ported in every kind of machine running a Java Virtual Machine (JVM). IBM is going to use open standard for each software product. IBM HTTP Server (powered by Apache) stems from Apache, an open source community that develops and maintains servers for modern systems. The October 2003 Netcraft Web Server Survey found that more than 64% of Web sites on the Internet are using Apache. See the Netcraft site for more information at:

http://news.netcraft.com/archives/2003/12/02/december_2003_web_server_survey.html

The IBM *autonomic computing initiatives* can reduce time for managing servers and resources, using tools to make technology self-managing. With IBM Tivoli System Automation for Linux, you have self-healing of business applications to detect failing IT components, and to cure or work around those failures. The products are available for Linux for zSeries and help you to address:

- ▶ High application availability
- ▶ Increasing complexity and operation costs
- ▶ Automation implementation and maintenance costs
- ▶ Rapid change of your IT infrastructure



xSeries and Linux

This chapter provides an overview of the Linux implementation, solutions, and resources available for the IBM @server xSeries server, which we will call just *xSeries* here for the sake of readability.

This chapter covers the following topics:

- ▶ How Linux runs on the xSeries
- ▶ xSeries customers and Linux
- ▶ xSeries and Linux as a server consolidation solution
- ▶ xSeries Linux customer
- ▶ Linux in a distributed enterprise
- ▶ Linux applications for xSeries
- ▶ Infrastructure solutions for xSeries
- ▶ xSeries architected solutions

5.1 X-Architecture™

The xSeries architecture principle is best understood once placed in its historical context, as the following table shows.

Table 5-1 The continuing evolution of needs 1960-2010

Decade	Urge of the customers	How IBM met these needs
1960's: The batch years	Binary compatibility between their computers at upgrade time.	The System/360
1970's: Time-sharing	Virtual memory to ease memory constraints	The System/370
1980's: The personal computer era	Robust, standard, modular and open personal computing platforms	The IBM PC
1990's: The Internet boom	RISC computers optimized for efficient scientific computations	The Risc System 6000 Series and its PowerPC component
2000's: Huge Web servers as commodity	24x365 availability of the servers, and hardware upgrades on-the-fly	Modular, scalable server technology, and pay-as-you grow: The IBM xSeries Servers.

5.1.1 Principles

The idea behind IBM X-Architecture technology is to extend the benefits of advanced availability technologies to IBM Intel processor-based servers. These benefits are in the areas of availability, scalability, systems management, service, and support.

Enterprise X-Architecture technology enables the following capabilities:

- ▶ XpandOnDemand™ scalability: You can add memory without having to stop the server.
- ▶ System partitioning (system partitioning is the process of dividing your system into non-overlapping sets of nodes to tailor your system to your needs).
- ▶ PCI-X I/O subsystem: Active™ PCI-X (PCI-X allows among other things a throughput of 1056 Mbytes/sec instead of the 528 Mbytes/sec of regular PCI).
- ▶ Remote I/O (remote I/O allows to locate I/O modules in bases at some remote distance from the CPU base. These remote bases are controlled by a remote master on the computer. Each remote base unit has a remote slave).
- ▶ Active Memory™
 - 256 GB memory capacity

- High-speed (DDR) memory
- Memory ProteXion™ (eliminates bad bits in the same way a disk drive provides transparent defective sector replacement)
- Chipkill™ memory (does for the RAM what RAID does for the disks)
- Memory mirroring
- Hot-add/hot-swap memory (yes, without stopping the server!)
- ▶ 533 MHz front side bus (FSB)
- ▶ Xcel4 Server Accelerator Cache
- ▶ Real-time diagnostics

An IBM xSeries server is designed for non-stop operations, 24x365, providing the operating system can take advantage of these functionalities. For this reason, IBM supplies appropriate Linux drivers for its products (such as ServeRAID™). Other drivers may be provided through redistribution or preload agreements.

IBM makes sure that its systems are well tested on Linux OS in its centers. Raleigh, for example, is a center where system hardware operating systems are tested. (Some of the hot add/swap functions described are supported only in certain versions.)

Today, IBM goes on building on the X-Architecture blueprint with Enterprise X-Architecture technologies. These technologies yield serious advances in the input/output (I/O), memory, and performance of xSeries servers. This new server design creates a flexible “pay as you grow” approach to buying high-end 32-bit and 64-bit xSeries servers. The results are systems that can be scaled quickly, easily, and inexpensively.

These features bring to industry-standard servers the kinds of capabilities that were formerly available only to users of mainframes and other high-end systems. Combined with existing X-Architecture technologies, these innovations result in large “economies of scalability,” excellent flexibility, and new levels of both server availability and performance for Intel processor-based servers.

In Intel-Architecture (IA) servers, Linux has evolved quickly. When compared with other operating systems, new drivers and functions were incorporated one after another from all the open source community, and put in the common pot.

From the day distributions utilizing kernel-2.4 appeared, Linux for Intel-Architecture servers has been fully accepted as an enterprise operating system. Other platforms do increase their number of servers, but the Intel-Architecture server market with Linux is now by far in the first place.

5.1.2 Context

IBM used its experience from years of enterprise server development to enhance the industry and used the industry-standard Intel CPU as a building block (Figure 5-1). The Intel processor is by itself inexpensive, and our engineers built a proper architecture around it to allow mainframe-quality server function. Concerning software, IBM worked on drivers so Linux could use these innovations in a way transparent for your IT Team and for applications.

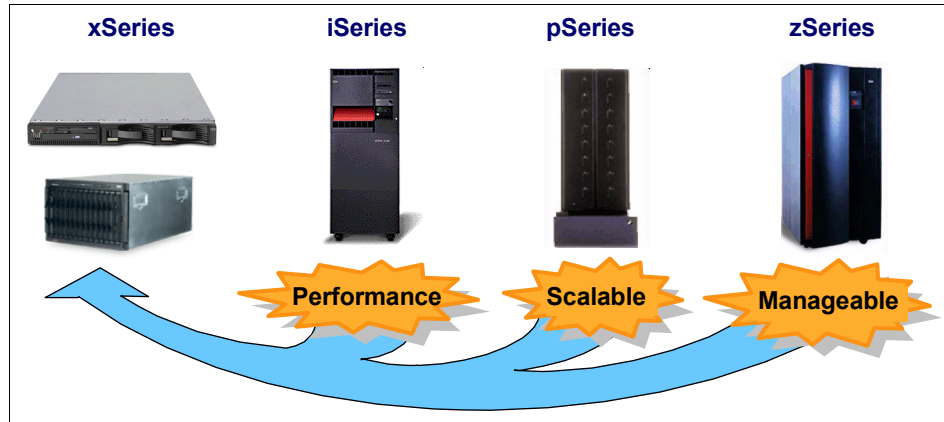


Figure 5-1 The concept of xSeries: Learning the best from all IBM @server product lines

While the xSeries server delivers advanced technology to customers at a low price, many applications and services with industry standards are also usable.

5.1.3 Models



Physically, IBM xSeries servers exist under the following forms, going further and further away from a basic PC for easiest integration of multiple systems in one place:

- ▶ Universal servers: the 200 series
- ▶ Rack-optimized servers: the 300 series
- ▶ High performance scalable servers: the 400 series
- ▶ IBM @serverBlade servers (Blades exist for both i386 and PowerPC architectures and they use the same racks)
- ▶ Some *pre-configured Linux clusters* (IBM @server1350 clusters) including among others a 4 x1U 335 servers in a 11U rack, a 32-node 335 models in 42U enterprise rack, etc.

Please be sure to check <http://www.ibm.com/servers/> in order to have an up-to-date list for both xSeries and prices. Technology evolves very quickly.

At the time of this writing, the xSeries is available in the following models implemented on the following models.

Table 5-2 Models of xSeries

Model	Type	Specificities
205	Universal Servers	Range from \$499 to \$4999 in December 2003, and from the 205 (tower with optional 4U rack capability, 128 MB to 2GB ECC DDR PC2100, hot-swap disks to 587 GB) to the 255 (7U chassis or tower, up to 4 Intel Xeon MP, up to 24 GB DDR ECC Chipkill, up to 12 hot-swap HDD bays supporting 1.76 TB of external storage)
225		
235		
255		
		Both IBM, through its ServerProven® Program, and Linux Distributors through their certification programs perform testing on all of these systems. To see the latest status of IBM ServerProven for a given system/OS combination, please visit: LINK To see the latest status for Linux Distribution Partner certifications for a given system/OS combination, please visit: LINK2
305	Rack-optimized	Range from the highly-affordable (\$1229 in December 2003) 305 server (1U chassis, Pentium 4 up to 3.06 GHz, 533 MHz FSB, 18 to 293 GB of SCSI disk - IDE possible - two 64-bit PCI-X slots) to the 382 (2U chassis, two Intel Itanium 2 64-bits processors, 2 to 16 GB of interleaved ECC DDR, three 64-bit PCI-X slots and integrated Dual Channel Ultra320 SCSI controller with RAID-1 for mirroring).
335		
345		
360		
365		
382		
		

Model	Type	Specificities
440	Pay-as-you-grow scalable servers	All these servers use the 4U form factor. They range from the 440 (upgradable up to 16 Xeon processors MP and 64 GB Chipkill memory per chassis) to the 455 (upgradable to 16 Intel Itanium 2 processors and 224 GB Chipkill memory per chassis). All of them offer 6 ActivePCI slots (hot plug) and are Remote I/O scalable (standard on some models, in option on other ones)
445		
450		
455		
HS20	Blade servers	Blade servers use one or two Intel Xeon (R) processors from 2.6 GHz to 3.06 GHz with 512KB cache memory; they range from 512 MB to 8GB with Error Correction (ECC), with handling of 146 GB of SCSI disks and 2 Gigabit Ethernet controllers per blade in December 2003. For up-to-date info at time of reading, please refer to: http://www.ibm.com/servers/ (some blades also use the Power PC processor; pick your own choice). Here is a blade:



Operating system support for xSeries changes frequently. Table 5-3 lists several helpful Web sites to which you can refer for platform-based information about the services that are available for the xSeries server. This information is available at the IBM ServerProven site on the Web at:

<http://www.pc.ibm.com/us/compat/nos/matrix.shtml>

You can find the drivers and documents for xSeries on the Web at:

http://www.pc.ibm.com/support?lang=en_US&page=brand&brand=IBM+PC+Server

At the time we write this page (December 2003), 83 drivers are available in category “Linux.”

Table 5-3 xSeries Linux information matrix

Information	Location on the Web
xSeries Linux information	http://www.pc.ibm.com/us/eserver/xseries/index.html
eServer BladeCenter™	http://www.ibm.com/servers/eserver/blades/
xSeries Linux clustering	http://www.ibm.com/servers/eserver/cluster
xSeries Linux independent software vendor (ISV) resources	http://www06.ibm.com/developerworks/offers/linux-speed-start/isv.html –OR– http://tinyurl.com/uno7
Linux Test Drive–xSeries	http://www-1.ibm.com/servers/enable/site/testdrive/xseries/ –OR– http://tinyurl.com/unq4

5.1.4 Linux and functions of X-Architecture

X-Architecture technology, the evolving blueprint for xSeries servers, considered much from the vast enterprise server heritage of IBM. xSeries engineers considered both the technologies that had already revolutionized larger IBM systems, and the best emerging technologies concepts. They applied this combination to the Intel-based platform, and now our customers have made the xSeries server the industry-standard. Getting enterprise power, scalability, control, and service at attractive prices was probably the reason of their choice. Table 5-2 lists the Linux support for X-Architecture.

Table 5-4 Linux support and X-Architecture

Functions	Supported by Linux
Active PCI (hot-add and hot-swap)	IBM-supplied drivers for Linux
Redundant NIC	
Advanced Systems Management processors	
IBM Director 4.1	IBM supports native Linux
Storage area network (SAN)	IBM-supplied drivers for Linux

Restriction: This supports may change depending on models you are evaluating. For more information about the entire platform supports and functions, refer to:

<http://www.pc.ibm.com/us/eserver/xseries/xarchitecture/>

You can download drivers from the Web through either URL:

5.1.5 Enterprise X-Architecture with Linux

Many customers want efficient support for both Microsoft and Linux operating systems, a key strength of the x-Architecture is designed to allow it. This enables customers to choose from an enlarged portfolio of applications to suit their business needs. To further enhance flexibility, physical partitioning and a common platform design help to ease the transition from 32-bit to 64-bit high-end computing.

Not only do new tools make systems management easier than before, but with self-diagnosing and self-healing technologies (essentially Active PCI-X and third-generation Chipkill memory), systems are designed to stay up and running continuously, even with software upgrades, and even with hardware upgrades. This availability combined with the ability to scale a server quickly, easily, and inexpensively makes it a platform of choice. Table 5-5 shows the Linux support that is available for Enterprise X-Architecture.

Table 5-5 Linux support with Enterprise X-Architecture

Functions	Correspondence by Linux
XpandOnDemand Scalability	Red Hat Enterprise Linux 3, SuSE SLES 8
Remote I/O	Kernel-2.4.9-31 or later

XpandOnDemand is part of the X-Architecture and will enable 16-way configuration on the xSeries 440 and 445 products. Powered by Enterprise X-Architecture™ technology, these 4U rack-optimized, industry-standard servers support up to 16-way processing by interconnecting two xSeries 440 chassis as a single 8U configuration. This makes is one of the most rack-dense 16-way servers in the world.

The xSeries 445 is the new level of high-performance enterprise server. This increase XpandOnDemand technology and is optimized for server consolidation and high-end database. The xSeries 445 combines scale-up or scale-out flexibility with the high performance of Intel Xeon. Table 5-6 shows the differences between the model x440 and x445.

Table 5-6 x440 vs. x445

x440	x445
XA-32 1st Generation Chipset	XA-32 2nd Generation Chipset
2-way to 16-way SMP	2-way to 16-way (4q03 = 32-way)

x440	x445
Intel Xeon MP: 1.5/1M, 1.9/1M, 2.0/2M	Intel Xeon MP: 2.0/1M, 2.5/1M, 2.8/2M
Intel Xeon DP 2.4GHz up to 4-way	Intel Xeon DP 3.0 GHz up to 4-way
32MB XceL4 per CEC. 128MB Max	64MB XceL4 per CEC, 256MB Max
32GB Max Memory per chassis	64GB Max Memory per chassis
PC133 SDRAM, 4-way Interleaving	DDR SDRAM, 2-way Interleaving
Adaptec Ultra160 SCSI, RAID Optional	LSI Ultra320 SCSI, Integrated RAID-1
Max Storage = 2 HDDs x 73GB = 146GB	Max Storage = 2 HDDs x 146GB = 292GB
Active PCI-X: 2 @ 133MHz, 2 @ 100MHz, 2 @ 66MHz	Active PCI-x: 2 @ 133MHz, 2 @ 100MHz, 2 @ 66MHz
Remote I/O Support	Remote I/O Support + RIO Sharing
24x CD-ROM UltraBay	8x DVD-ROM UltraBay
Broadcom 5700 single port GbE	Broadcom 5704 dual port GbE
Active Memory = ECC, Chipkill, Memory Mirroring & Memory ProteXion	Active Memory + Hot-swap Memory and Hot-add Memory
1-year Same Business Day 24x7 Warranty	3-year Next Business Day 9x5 Warranty
4U: 19"(483mm) x 7"(178mm) x 27.5"(698mm)	4U: 19"(483mm) x 7"(178mm) x 28.1"(714mm)
35 #1 Benchmarks in 15 months	Up to 45% Performance Improvement*

*As measured by TPC-H comparing x440 to x445

XpandOnDemand offers scalability in the way you buy and grow. This revolutionizes data center servers with a modular, pay-as-you-grow building block design. This design offers low entry price points and upgradability to powerful 16-way SMP and remote I/O. XpandOnDemand allows you to purchase only the performance and I/O capacity that you need, when you need it, without having to buy costly upfront infrastructure.

For more information about Enterprise X-Architecture, see:

<http://www.pc.ibm.com/us/eserver/xseries/xarchitecture/enterprise/index.html>

Note: In order to be able to use the Hyper-Threading Technology, and the NUMA features of x440/x445, you must use a 2.6 kernel-based distribution or substitute a 2.6 kernel to the 2.4 in an existing distribution.

5.1.6 xSeries Linux distributions

Linux for the xSeries server is available from leading Linux distributors: Red Hat, SuSE, and Turbolinux. Recently Conectiva, SuSE, and Turbolinux partnered to create UnitedLinux. This new initiative may streamline Linux development and certification around a global, uniform distribution of Linux, designed for business. IBM will continue to support Red Hat Linux and fully support UnitedLinux, which should make it easier to create a wide variety of Linux solutions.

For a detailed explanation of each distribution, see 1.9, “The Linux distributions” on page 24. You can also refer to the following Web sites:

- ▶ Red Hat Linux
<http://www.redhat.com>
- ▶ SuSE Linux
<http://www.suse.com>
- ▶ Turbolinux
<http://www.turbolinux.com>

The packages of these distributions are updated frequently. To download fixes and security packages, see the following Web sites:

- ▶ Red Hat
<http://www.redhat.com/apps/support/errata/>
- ▶ SuSE
<http://www.suse.com/us/private/download/updates/index.html>
- ▶ Turbolinux
<http://www.turbolinux.com/security/>

5.1.7 The Red Hat: IBM announcement of December 2003

IBM and Red Hat jointly announced in December 2003 the availability of the just-released Red Hat Enterprise Linux 3 (RHEL3) across the whole IBM product line, and supporting WebSphere, DB2, Lotus, and Tivoli as well, with a total of more than 200 software products. The two companies will provide one-stop support for customers' Linux requirements.

More details are available at: <http://www.ibm.com/linux/> and <http://www.redhat.com/software/rhel/>

5.2 Intel-based server customers and Linux

Analysts estimate that the market for Intel-based servers (IA-32 and IA-64) will grow at about 10% compound annual growth rate (CAGR) between now and 2005. They see the 64-bit market for Intel servers showing significant growth. In the expanded market for Intel-based servers, Linux shows the fastest growth of all operating systems.

Most of the current Linux usage comes from small businesses. IBM market research shows that most of the deployment of Linux applications is in the infrastructure server area. Within two years, the share of more important applications such as the database server will expand significantly. Table 5-7 shows several motivations to move to a Linux platform based on IBM market research from February 2001. There is an increasing interest in Linux for line-of-business (LOB) use in large enterprises because of its increasing availability, reliability, and customer successes.

Table 5-7 Worldwide Linux application deployment

Application deployment	Today	Within two years
Web server	51%	76%
Web application server	34%	61%
Firewall server	35%	60%
E-mail server	40%	59%
Network server	39%	59%
Database server	29%	55%
Development system	29%	55%
Workgroup server	18%	40%
Transaction server	11%	22%

The most influential factors in moving to Linux can reduce overall cost. Reliability is proving to be a significant factor. As Linux moves into the mainstream, industry acceptance will increase as indicated by the chart in Figure 5-2.

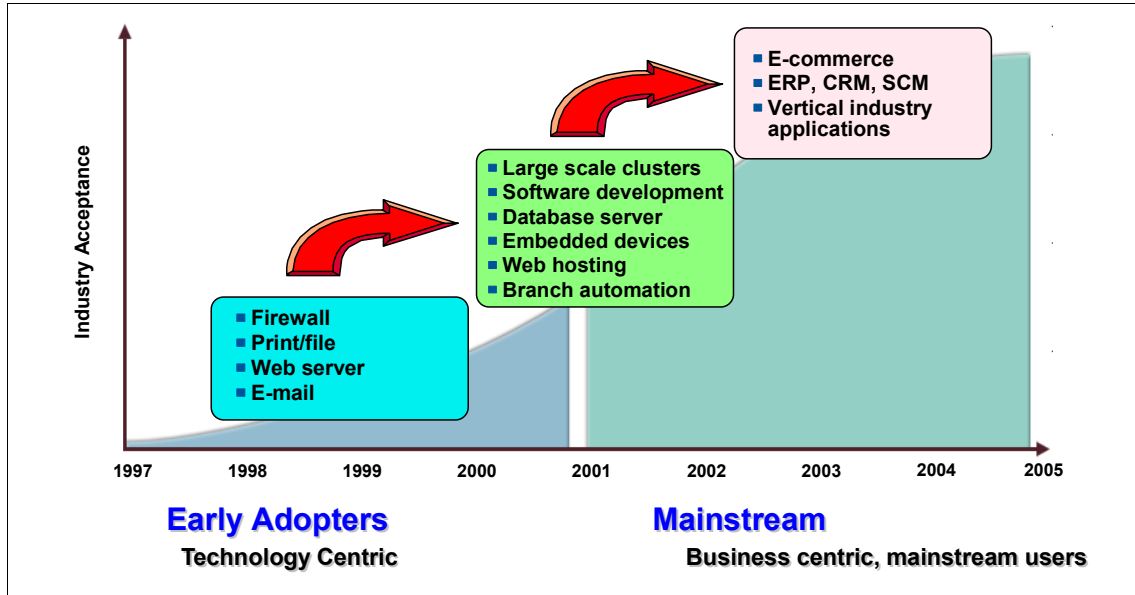


Figure 5-2 Expected growth rate of Linux acceptance

5.3 Consolidation

There are four types of workload consolidation (also referred to as *server consolidation*). Each type offers significant benefits in the following areas:

- ▶ Reduced administrative costs because of central management
- ▶ Better management of system proliferation and more consistent architecture
- ▶ Management of purchasing to achieve volume purchasing discounts
- ▶ Consistent process for security, operating system levels, and updates

There are four types of consolidation:

- ▶ **Consolidating multiple systems previously spread around the enterprise into fewer, more centralized locations**

For this type of consolidation, xSeries and Linux together provide the following benefits:

- Reduced floor space: The xSeries features rack-dense solutions and cable-chaining technology.
- IBM Director: This includes a function called Update Xpress, which allows you to propagate system images and updates to all consolidated systems.

- Workload management solutions: They reschedule work from heavily loaded systems to more lightly loaded systems.
 - Availability through redundancy within systems and across systems: With a consolidated environment, you can configure your solutions for maximum availability. The xSeries server already has redundancy within systems for such things as power supplies. Now, in a consolidated environment, you can set up another system to handle work from a failing system or from a system you may want to take offline to replace.
- **Consolidating many systems into a reduced number of larger servers**
- The xSeries 400 series offers a strong value proposition with its *Pay as You Grow* feature. This feature allows you to buy only the computer resources you need for the first stage of your consolidation efforts, and to add additional computer capacity as your consolidation efforts grow or business expands.
 - The remote I/O feature of the xSeries models 400 allows you to increase the I/O capacity of your systems without throwing out the base system. You can expand your computer power with “Pay as You Grow” and customize the I/O needs as your consolidation requirements change.
 - Virtual or physical partitioning: The xSeries servers allows you to create physical partitions, each with its own memory and I/O. In addition, VMware offers virtual partitions. With both partitioning configurations, failover between virtual or physical partitions is possible.
 - This environment requires fewer administrators.
 - This environment offers separate resource allocation, billing, and server-level agreements.
 - VMware and virtual partitions: VMware offers several key benefits as you implement server consolidation. These benefits include:
 - The ability to run Windows and Linux partitions side by side. You can run Windows applications in one partition, and Linux applications in a second partition. VMware supports both SuSE and Red Hat as Linux operating environments.
 - The ability to change partitions dynamically as computer needs change. You can add a new partition or alter the I/O resources dedicated to it in a few minutes
 - You can set up partitions and customize them to serve specific workload requirements. For example, you can set up two Web-serving partitions, a print-serving partition, and two mail-serving partitions. The mail and Web-serving partitions may likely have mission-critical requirements, so you can set up the second partition as a backup partition, which you can shift workload to if one partition has capacity

problems. The print-serving partitions are likely set up with less I/O capacity dedicated to them than to the mail-serving partition. This makes customization very flexible for your specific workloads.

► **Consolidating data**

Main benefits that xSeries and Linux together bring to this area include:

- The ability to create highly available data configurations with system area networks and shared storage environments
- Data mirroring to ensure availability of data if disk failures occur
- Improved security with centralized security policies rather than fragmented data scattered around the enterprise
- Fewer database administrators required

► **Consolidating applications**

The main benefits that xSeries and Linux together bring to this area include those mentioned above as well as these:

- A reduction in application licensing costs by consolidating application usage on one system rather than multiple systems across the enterprise
- The ability to better manage new application migration plans with operating system or application-level changes
- A reduction in compute capacity requirements by load balancing multiple application requirements across systems

VMware ESX Server

With the introduction of VMware ESX Server, the xSeries server now offers virtual partition functionality that was previously only found in main-frame environments. VMware ESX Server simplifies server infrastructure by partitioning and isolating server resources in secure and transportable virtual machines. VMware ESX Server enables you to remotely manage, automatically provision, and standardize these server resources on a uniform platform.

Figure 5-3 shows the architecture of VMware ESX Server. Operating systems and applications are isolated in separate virtual machines.

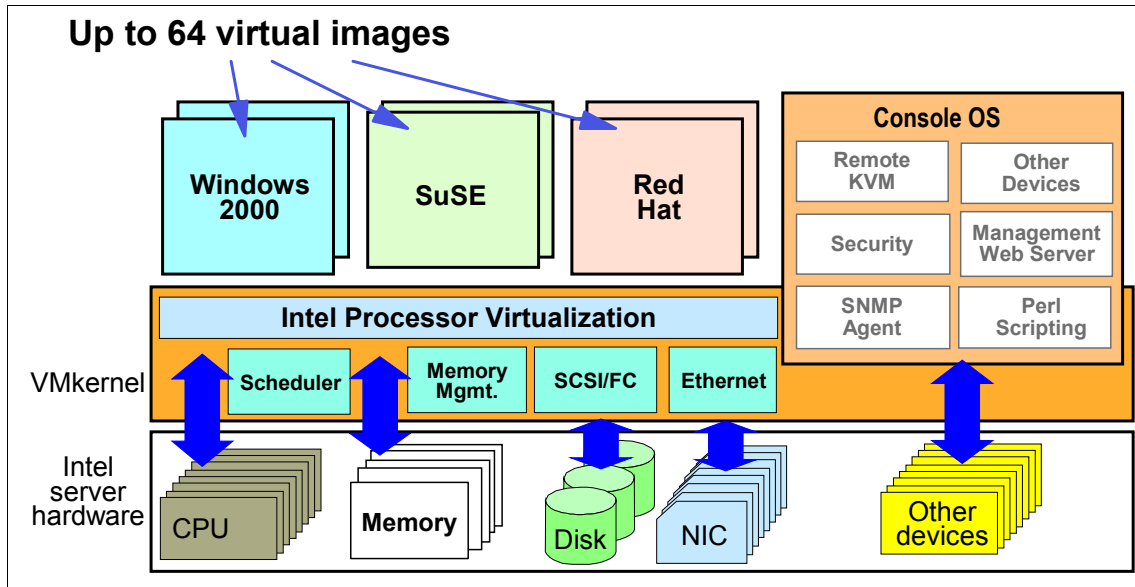


Figure 5-3 The architecture of VMware ESX Server

VMware ESX Server transforms physical systems into a pool of logical multiple virtual machines that reside on a single piece of hardware. System resources are dynamically allocated to any operating system based on immediate need.

VMware ESX Server runs directly on system hardware to provide a secure, uniform platform for easily deploying, managing, and remotely controlling more operating systems. Advanced resource management controls allow you to guarantee service levels of CPU, memory, networking, disk resources, and failover between partitions.

VMware ESX Server offers the following benefits:

- ▶ You can move applications running on dedicated systems into separate virtual machines on a single, more reliable, and scalable system.
- ▶ You can manage servers remotely from any location, simplifying server maintenance.
- ▶ With advanced resource management controls, you can guarantee service levels.
- ▶ You can script common monitoring and management tasks.

For more information about VMware ESX Server, see this site:

http://www.vmware.com/products/server/esx_features.html

5.4 Linux clusters

A clustered system is basically a system that combines two or more computers into a group to improve reliability or processing performance. The types of clusters can be roughly classified as:

- ▶ High-availability cluster
- ▶ Load-balancing cluster
- ▶ High-performance computing
- ▶ Tightly coupled clusters

For a detailed explanation and references about clusters, see Chapter 10, “Linux clusters” on page 313.

5.4.1 High-availability cluster

As your computer requirements expand, you can choose to expand your systems vertically or horizontally. *Vertical scaling* essentially means implementing larger systems. *Horizontal scaling*, or clusters, means implementing multiple smaller systems to handle the workload and leveraging them either as one computer resource or as a large server farm with sophisticated workload balancing.

Cluster implementations begin with high-performance computing environments where customers have hooked many systems or nodes together over a high-speed interconnection. By leveraging software, they split jobs into small chunks to increase speed and efficiency.

Today, clusters are implemented in life sciences, financial, and petroleum environments, where large amounts of data are analyzed and the problems are well suited to be split among multiple machines. Often, clusters are implemented as large server farms, an implementation that has many of the benefits of server consolidation.

IBM is a leader in Intel-processor Linux clusters. The key differentiators include:

- ▶ **Reduced system cost:** With Linux and xSeries solutions, the initial cost is extremely attractive.
- ▶ **System management:** With IBM Director and Cluster System Management, you now have solutions to manage many nodes as one computer resource. This produces savings in total cost of ownership (TCO) by requiring fewer administrators and improving system availability.
- ▶ **High-availability solutions:** IBM, together with its ISV partners, offers high-availability solutions to allow failover from failing systems within the cluster to nodes that are up and operational.

- ▶ **Cluster integration services:** Clusters are difficult to implement and design. The strategy must include application and workload analysis. It must also include a plan for how all the pieces are put together, and how well all the pieces come together at installation time. Cluster integration services provided by an IBM Business Partner or IBM Global Services (IGS) are a key differentiating value proposition that has produced many successful cluster implementations.

High-availability software

The software for high-availability clusters includes the following products. Consult the accompanying Web site for more information:

- ▶ Veritas Cluster Server
<http://www.veritas.com>
- ▶ SteelEye Lifekeeper
<http://www.steeleye.com/products/linux/>
- ▶ Toshiba DNCWARE Cluster Perfect
<http://www.toshiba.co.jp/index.htm>
- ▶ Red Hat Cluster Manager
<http://www.redhat.com/software/rhel/as/cluster/>
- ▶ Linux Virtual Server
<http://www.linuxvirtualserver.org/>
- ▶ Tivoli System Automation
<http://www.ibm.com/software/tivoli/products/sys-auto-linux/>

5.4.2 Load-balancing cluster

Load-balancing clusters raise performance and throughput by distributing many jobs, processes, and requests to two or more servers. Load-balancing cluster software has a flexible setup process, which is a strong point.

Advanced load balancing has an agent that is installed on each nodes. This client performs the health check of the database server, the application server, etc.

Load-balancing software

The load-balancing cluster applications include three major software providers:

- ▶ IBM WebSphere Edge Server
<http://www.ibm.com/software/webservers/edgeserver/index.html>

- ▶ Turbo Linux Cluster Server
<http://www.turbolinux.com/products/tcs/>
- ▶ Ultra Monkey
<http://www.ultramoney.org/>

5.4.3 High-performance computing

High-performance computing performs calculations that are numerically intense and require high performance and a mass memory. The tasks are carried out by parallel execution of the independent processes.

There are two types of high-performance computing: one that uses a library for message passing and the other that uses a job scheduler. When the library is used for message passing, an independent process is performed in two or more threads and processes.

When the job scheduler is used, one application is simultaneously run in two or more nodes. Each node performs the process with a script or an exclusive tool.

IBM Cluster 1350

Cluster 1350 is the latest generation of affordable integrated Linux cluster offerings from IBM. Cluster 1350 combines the power of the latest xSeries rack-optimized servers with IBM Cluster Systems Management (CSM) for Linux software, IBM storage products, and leading third-party networking components. It offers powerful, flexible solutions for high-performance computing (HPC) and commercial application environments.

Cluster 1350 is available in a wide range of configurations, including single-rack systems of up to 32 cluster nodes and large-scale systems of up to 512 cluster nodes and beyond. Cluster nodes are one- or two-way x335 servers (1U) with Intel Xeon processors running at 2.4, 2.6, or 2.8 GHz. Systems may optionally include up to 32 one- or two-way x345 storage servers (2U), with Intel Xeon processors running at 2.4, 2.6, or 2.8 GHz. In addition, the xSeries Model 345 can be used as a cluster node. Each cluster requires a two-way Model 345 management node.

Cluster 1350 systems are managed by IBM Cluster Systems Management for Linux Version 1.3. This provides resource monitoring, automated operations, remote hardware control and command execution, configuration file management, and parallel network installation. CSM for Linux helps ease administration and may reduce life-cycle costs by allowing management of an entire Cluster 1350 system from a single management node. As workload demand increases, CSM for Linux allows incremental growth of the Cluster 1350 configuration without necessarily increasing the management complexity.

The highlights of Cluster 1350 include:

- ▶ Hardware
 - Rack-optimized xSeries 335, 345 and 360 Intel processor-based
 - Rack-optimized xSeries 325 AMD Opteron-based and IBM BladeCenter
 - Industry-standard 10/100Mbps Ethernet or Gigabit Ethernet cluster interconnect
 - Optional high capacity IBM FASTT storage
 - Terminal server and KVM switch
 - Hardware installed and integrated in high density 42U enterprise racks
- ▶ Operating system
 - SLES 8/AMD (64-bit)
 - RHEL 3.0 AS/AMD and WS/AMD (64-bit) (available 1/04)
 - RH 9 (32-bit only)
 - RHEL AS 2.1
 - SLES 8
 - RH 9
 - RHEL 3.0 ES and WS (available 1/04)
- ▶ Cluster software
 - IBM Cluster Systems Management (CSM) for Linux 1.3.2
 - IBM General Parallel File System (GPFS) for Linux 1.3.0-3 (optional)
 - Linux CLuster Install Tool (LCIT)
 - Scali MPI Connect
- ▶ Services
 - Base system integration and setup
 - Warranty support
 - Linux Cluster Installation Services
 - Support Line for Linux Cluster

Table 5-8 shows a sample configuration available for Cluster 1350.

Table 5-8 Cluster 1350 Sample configuration

Level	Cluster size	Processor speed	System memory	Internal storage	Cluster interconnect
Entry	8 cluster nodes	2.4 GHz	512 MB	18 GB	10/100Mbps Ethernet
Mid-range	32 cluster nodes	2.8 GHz	1 GB	18 GB	10/100Mbps Ethernet
High-end	128 cluster nodes	3.06 GHz	1 GB	36 GB	Gigabit Ethernet

Level	Cluster size	Processor speed	System memory	Internal storage	Cluster interconnect
High-performance	64 cluster nodes	3.06GHz	1 GB	36 GB	Myrinet-2000

For more information about Cluster 1350, see:

<http://www.ibm.com/servers/eserver/clusters/>

5.5 Linux in a distributed enterprise

Complex enterprises such as retail stores and banks often require computers in many sites. The applications run on many geographically dispersed machines and are managed remotely. Because of its low cost and small footprint, Linux on the xSeries is an excellent solution. In distributed enterprises, ease of use, and ease of support are critical. Once again, the xSeries and Linux are perfect for this situation because of their high availability and reliability. Linux is also a great solution as a standard application interface that links all applications on the xSeries servers.

For example, Linux provides a large retail outlet with functionality that is the same as or better than Windows NT at about 25% lower TCO, including hardware, software, operating system licensing, and support costs.

Note: At the time the retail outlet implemented the solution, Sun was charging for the Solaris operating system. However, Sun no longer charges a licensing fee for Solaris for low-end hardware.

Linux provides a standardized interface for the deployment of applications, enabling a reduction in deployment time. With networked terminals, administrators can upgrade applications once per site, rather than once per computer. Linux provides the retail outlet with a standardized platform for the deployment of new applications, reducing application deployment time by about 50%. With retail outlet's Linux solution, application upgrades are about 400% faster with the thin client architecture because programs run remotely, and administrators are not required to load as many applications.

Distributed enterprise can reduce support requirements by maintaining one copy of an application at the server level, rather than maintaining multiple copies at the client level. In some cases, support time and personnel required may decrease by up to 90%. Aggregated over many stores or locations, some customers support savings can reach millions of dollars.

5.5.1 Basic implementation strategies

There are three basic implementation strategies for distributed enterprises:

- ▶ **Decentralization**
 - Relocates information technology (IT) operations closer to the user
 - Delivers to the user from a local server that communicates to the user through a thin client
- ▶ **Centralization of corporate data assets**
 - Provides significantly improved business control capabilities and rapid rollout and deployment of enhanced function
- ▶ **Application simplification**
 - Provides identical application functionality to each user through thin client or kiosk
 - Reduces the management and maintenance of servers and applications
 - Reduces TCO

5.5.2 Business benefits

The distributed enterprise solution provides four key business benefits:

- ▶ Centralized management through the server and Internet for a secure, stable desktop environment
- ▶ Accelerated deployment of applications, capability, hardware for speed and flexibility
- ▶ Access to new and existing applications implies seamless access to multiple servers and networks
- ▶ Low TCO
 - Centralized management; lower training and support cost
 - Rapid application and function deployment
 - Significantly reduced client functionality requirements
 - Reduced software licensing costs

5.5.3 Customer example: A large online company for trading stocks, bonds and other personal investment activities

This online personal investment company knows what it takes to become a leading financial services brand. With the launch of their Web site in 1996, the company experienced an immediate explosion in consumer demand for its services. From its beginnings as a discount brokerage firm, the company has

grown into a global leader in personal financial services with branded Web sites spanning the globe and total customer assets.

When planning an upgrade to its Web services, the company decided to replace the hardware and software that powers its North American financial services business. After careful analysis, the company replaced its servers with IBM xSeries servers. The company also made a fundamental change in its operating system. It resolved to deploy Linux, because it is an open-source and standards-based platform and holds the promise of cost savings.

With 10,000 concurrent users and four million customer accounts, the company's success depends on its ability to provide superior customer service. As a fully integrated financial services firm with a significant online presence, providing a high-performance Web site that is up and running whenever customers want to conduct business is the essence of the brand. The company must also have a very scalable system that can handle the peaks and valleys of the online financial services business.

The company recognized the potential of Linux, and needed a complete solution from a vendor who also recognized and embraced Linux technology, and had experience with large Linux deployments. In switching to Linux, the company was positioned in the vanguard of the open-source movement and as a financial services technology leader. To make the transition, the company needed a trusted provider. They chose IBM.

The company upgraded its data center operations by deploying 90 xSeries Model 330 servers running Linux. By using Linux, the company sharply reduced software licensing fees and service costs. The company's new deployment is expected to provide the savings and performance levels necessary to meet business goals. In addition, because the solution is based on open standards, the company has flexibility in the applications the company chooses to deploy.

The company believes that the combination of the xSeries servers and the IBM commitment to Linux will save them millions of dollars and enhance the the company's global brand. By using Linux, the company realized sizable software cost savings. By deploying the lower-cost xSeries servers, the company has experienced a rapid return on investment.

For more information about this success story, see either URL:

<http://www.ibm.com/software/success/cssdb.nsf/CS/KMOE-5EAMKE?OpenDocument&Site=linuxatibm>

5.6 Linux application solutions

Although the general opinion is that Linux has a limited number of enterprise-level applications available, the Linux applications featured in the IBM Global Solutions Directory have grown to more than 2,900. IBM has 10 porting centers around the world and is continually enhancing this portfolio.

5.6.1 IBM Software for Linux on xSeries

Just about every IBM software component needed to help customers enable their e-business is available for Linux including WebSphere, DB2 Universal Database, Tivoli, and Lotus Domino. The following sections list the readiness of IBM Software for Linux.

Note: The distribution level that IBM Software requires changes frequently. For more information, see Chapter 9, “IBM Software for Linux” on page 265.

WebSphere software platform

WebSphere is infrastructure software for dynamic e-business, delivering a proven, secure, and reliable software portfolio. WebSphere Application Server, MQSeries, and other product groups and development tools allow the latest Web technology, and enable simple extensions of e-business to meet market needs.

For more information about the WebSphere family, see this Web site:

<http://www.ibm.com/software/info1/websphere/index.jsp>

DB2 product family

DB2 Universal Database is used as the basis of e-business. It is a powerful database that provides the small and middle-scale user flexibility in response to demand, and provides the large-scale user the optimum relational database. DB2 supports industry-standard open-source software such as Linux, Java, and XML.

For more information about the DB2 family, see this Web site:

<http://www.ibm.com/software/data/>

Tivoli software product

Broadband provides a common connection for mass communication to the Internet, and e-business expands this connection even further. Many companies are required to provide systems 365 days-a-year, 24 hours-per-day. To provide this service requires an integrated management of the e-business infrastructure.

Tivoli can provide this integrated system management software. The management solution offers flexibility for the needs of e-business such as reduction of management cost, improvement in reliability of the whole system, and efficient management of the expanding data set. The open architecture, which corresponds to the multiplatform support, and the high scalability, which supports growth and change of business, are key features of the Tivoli product.

For more information about the Tivoli family, follow this link:

<http://www.ibm.com/tivoli>

Lotus Domino

The Domino family provides a multiplatform foundation for collaboration and e-business, driving solutions from corporate messaging to Web-based transactions.

For more information about the Lotus Domino family, see their Web site:

<http://www.lotus.com/>

5.6.2 ISVs and application solutions

The list of ISV products that are supported under Linux continues to grow significantly. Many major software vendors have made a commitment to the Linux platform. The ServerProven program from IBM ensures real business applications and solutions that are tested for the marketplace. To find the right solution, check out the growing list of validated ServerProven applications for Linux on xSeries.

For more information about ServerProven, visit this Web site:

<http://www.developer.ibm.com/welcome/eserver/eSC.pl?mvcid=Welcome&packageid=000>

Also, check the interactive Solution Sizing Tool Web site. It provides the best-fit server solution required to run a particular application environment. Refer to either of the following sites:

<http://www.developer.ibm.com/welcome/myvc.pl>

<http://myvctest.austin.ibm.com/servlet/>

[ServletEngine?mvcid=Welcome&packageid=000](#) –OR–

<http://tinyurl.com/unue>

The sizing tools that are available for Linux solutions include:

- ▶ **Apache HTTP Server for Linux:** The xSeries Apache 1.3 on Linux solution sizing tool is based on recommendations of xSeries servers running Apache 1.3 on Linux, and optionally SteelEye's LifeKeeper high availability clustering.

- ▶ **BISON Solution Release A2.9:** The xSeries BISON Solution Release A2.9 solution sizing tool is based on recommendations of xSeries servers running BISON Solution Release A2.9
- ▶ **Changing Worlds ClixSmart Navigator Server:** The xSeries Changing Worlds ClixSmart Navigator Server solution sizing tool is based on performance measurements of xSeries servers running Changing Worlds ClixSmart Navigator Server.
- ▶ **Check Point VPN-1 Pro:** The xSeries Check Point VPN-1 Pro solution sizing tool is based on performance measurements of xSeries servers running Check Point VPN-1 Pro.
- ▶ **DB2 UDB EE 7.2:** The xSeries DB2 UDB EE 7.2 solution sizing tool is based on recommendations of xSeries servers running DB2 UDB EE 7.2.
- ▶ **IBM HTTP Server for Linux:** The xSeries IBM HTTP Server for Linux solution sizing tool is based on recommendations of xSeries servers running IBM HTTP Server for Linux.
- ▶ **Integrated Sensors RTEExpress:** The xSeries and Netfinity® Integrated Sensors RTEExpress solution sizing tool is based on recommendations of xSeries and Netfinity servers running Integrated Sensors RTEExpress.
- ▶ **Lecando Training Server 5:** The xSeries Lecando Training Server 5 solution sizing tool is based on performance measurements of xSeries servers running Lecando Training Server 5.
- ▶ **Linux Gateway (Firewall & Proxy):** IBM xSeries Linux Gateway (Firewall and Proxy) solution sizing tool is based on the Linux Kernel 2.4.5.
- ▶ **mip Webbased Business Solution:** The xSeries mip Webbased Business Solution solution sizing tool is based on performance measurements of xSeries servers running mip Webbased Business Solution.
- ▶ **Net-Tech Linux ASP Solutions:** The xSeries and Netfinity Net-Tech Linux ASP Solutions recommendations are based on performance measurements of xSeries and Netfinity servers running Net-Tech Linux ASP Solutions by Net-Tech Linux Limited.
- ▶ **RealServer 8.01:** The xSeries and Netfinity RealServer 8.01 recommendations are based on performance measurements of xSeries and Netfinity servers provided by RealNetworks.
- ▶ **Samba File and Print Server:** The xSeries Samba File and Print Server solution sizing tool is based on recommendations of xSeries servers running Samba and optionally SteelEye's LifeKeeper high availability clustering.
- ▶ **Sendmail Mail Server:** The IBM xSeries Sendmail Mail Server solution sizing tool is based on recommendations of xSeries servers running Sendmail Mail Server and optionally SteelEye's Lifekeeper high availability clustering.

- ▶ **Shuttlesoft Volanto.Office:** The xSeries and Netfinity Shuttlesoft Volanto.Office recommendations are based on performance measurements of xSeries and Netfinity servers provided by Shuttlesoft.
- ▶ **Srishti PayScribe 3.1 for Linux Quick sizing tool:** The xSeries and Netfinity Srishti PayScribe 3.1 for Linux recommendations are based on performance measurements of xSeries and Netfinity servers provided by Srishti.
- ▶ **Trustix Firewall:** The xSeries Trustix Firewall solution sizing tool is based on recommendations of xSeries servers running Trustix Firewall.
- ▶ **Trustix Mail Server:** The xSeries Trustix Mail Server solution sizing tool recommendations are based on performance measurements of xSeries servers running Trustix Mail Server.
- ▶ **WebSphere Application Server for Linux:** The xSeries IBM WebSphere Application Server for Linux solution sizing tool is based on recommendations of xSeries servers running IBM WebSphere Application Server for Linux.

The following sections list examples of ISV applications that have been ported to Linux.

Non-IBM middleware

Today, Linux is used as an Internet infrastructure server and to build e-business environments. As important middleware, Oracle and Weblogic are mentioned as are IBM products:

- ▶ **Oracle9i Database** is a popular object-oriented relational database. It provides scalability, reliability, and security with rich functions. Oracle9i also has a clustering option named Real Application Clusters (RAC) that provides more scalability, availability, and performance for database applications. For more information, see:
<http://www.oracle.com/ip/deploy/database/oracle9i/>
- ▶ **Oracle9i Application Server (Oracle9iAS)** has an integrated J2EE application server. It provides cost-effective built-in portal function, wireless and voice, Web page caching, powerful business intelligence features, complete integration, and more, integrated in a single product. Oracle9iAS supports major J2EE, Web services, and XML industry standards. For more information, see:
<http://www.oracle.com/ip/deploy/ias/>
- ▶ **BEA Weblogic Server** is a Web application server that offers scalability, pliability, and reliability suitable for being called an enterprise application platform. For more information, see:
<http://www.bea.com/products/weblogic/server/index.shtml>

- ▶ **VMware ESX Server** simplifies server infrastructure by partitioning and isolating server resources in secure and transportable virtual machines. It enables these server resources to be remotely managed, automatically provisioned, and standardized on a uniform platform. For more information, see:

http://www.vmware.com/products/server/esx_features.html

IBM offers outstanding service and support for this product. IBM is the only hardware vendor that offers complete support for Windows, Linux, and key applications all in a VMware environment. VMware provides support for the VMware products only.

The process to obtain this high level of support is:

- a. **Register for Subscription Software Services:** One year of maintenance subscription service is included with your purchase of ESX Server from IBM. This entitles you to any updates or fixes for the VMware ESX Server software released within a year of your purchase, direct from VMware. To receive these updates, you must first register for this service with VMware. You must register for Subscription Software Services within 30 days after purchase.
- b. **Purchase Support for ESX Server:** Support for ESX Server is available through IBM Operational Support Services – Support Line. A single, integrated remote support solution, Support Line helps support your complex, distributed IBM and multi-vendor operating systems and software products. Contact IBM Global Services for more information on the Support Line offering in your area.

Enterprise applications

In recent years, the use of Linux enterprise applications has increased due to the stability of Linux. In the Linux enterprise application arena, companies such as SAP and J.D.Edwards provide solutions. Consider the following solutions:

- ▶ **SAP** is a comprehensive software package that manages the flow of the business processes of a company as represented by R/3. In the ERP software package market, SAP's market share reaches 50% or more, and serves as the global standard. In an Intel-Architecture server, since large-scale storage systems or system management software is not perfect, its use in core business is not very prevalent. However, Linux is expected to move into this area in the future. For more information, see:

<http://www.sap.com/linux/>

- ▶ **INSUITE Enterprise** is a Web-type groupware that offers functions for using information quickly and exactly in an all-in-one solution. It can build groupware, EIP, and a knowledge management environment in a short time. INSUITE is usable across the business from day one.

- **Cybozu Garoon** is a key collaboration application. Workgroup members can check mail, schedules, intranet documents, and various news forums through the Internet by using Cybozu Garoon from any Web browser. For more information, see:

<http://cybozu.com/index.cbml>

Security

The Linux server is used as a gateway server or file server in many cases in an Internet infrastructure. Although there are not many reports of virus infection in Linux environments at present, when a transaction that passes through a Linux gateway contains a virus, and the data resides in a Linux file server, the Windows system machine using the data can suffer damage.

The xSeries server offers a low-cost, secure foundation for your investment in Linux infrastructure and business applications. By providing hardware security features, self-healing, and self-diagnostic services, as well as secure remote administration and monitoring, the xSeries server is the platform for those who need to implement a secure system infrastructure.

IBM works closely with the leading security software providers and the managed security providers to deliver tested and validated security solution configurations and ServerProven security software. IBM security services augment the xSeries features by enabling customers to assess their security risks and develop plans to address the risks. The combination of a secure platform, a full line of solutions, and complete services coverage makes the xSeries server solution an ideal secure server solution.

Server security can be provided at several levels, including hardware, operating system, middleware, and applications. Table 5-9 summarizes how xSeries delivers security at each level.

Table 5-9 Levels of server security

	Security features	Customer value
xSeries hardware	<ul style="list-style-type: none"> ► Secure remote monitoring ► Secure remote administration ► Configuration control ► Lights-out support ► Self-diagnostics ► Hardware security framework 	<ul style="list-style-type: none"> ► Low-cost, secure platform for any security application environment ► Part of a complete IBM security solution built on IBM and industry-partner offerings

	Security features	Customer value
Operating system	<ul style="list-style-type: none"> ▶ Linux distribution ▶ OpenSSL encryption ▶ OpenSSH authentication ▶ PAM authentication ▶ SNORT intrusion detection ▶ Apache secure Web server ▶ Integrated firewalls in many distributions ▶ Tripwire for file integrity ▶ SuSE Linux Enterprise <ul style="list-style-type: none"> – OpenSSH – OpenSSL – PAM – Kerberos authentication ▶ Red Hat Advanced Server <ul style="list-style-type: none"> – Firewall – OpenSSH – OpenSSL – Kerberos authentication 	Linux is a low-cost, highly functional, and secure UNIX derivative. It brings the benefits of UNIX at a substantially lower cost. A complete suite of open-source and industry security solutions make Linux an enterprise-class environment. Linux is fully supported by IBM security services.
Middleware	<ul style="list-style-type: none"> ▶ Tivoli Risk Manager ▶ Tivoli Access Manager ▶ Tivoli Identity Manager 	A complete security management environment that can integrate with other IBM and ISV offerings and is fully supported with IBM services.
Applications	Trustix	<ul style="list-style-type: none"> ▶ Mail server with server-side antivirus ▶ File/Print
Security applications	<ul style="list-style-type: none"> ▶ Checkpoint ▶ Trustix ▶ Trend Micro ▶ Tripwire Security Systems, Inc. ▶ SuSE Linux firewall 	<ul style="list-style-type: none"> ▶ Firewall ▶ Secure Linux operating system ▶ Firewall with VPN ▶ Proxy server ▶ Antivirus software ▶ File-integrity security and policy-compliance applications
Services	<ul style="list-style-type: none"> ▶ IBM Security Assessment Services ▶ IBM Managed Security Services ▶ IBM Detection Services ▶ IBM Security and Privacy Services ▶ IBM Tivoli Consulting Services ▶ IBM Cross-Industry & Industry Solutions 	A comprehensive approach to security from a people and IT perspective. Addresses process, education, training, vulnerability assessment, and response plans. Helps customers identify critical assets and strengthen measures to control their access and use. The customer chooses whether to be security self-sufficient or to use IBM-managed services to protect their corporate and IT assets.

Here are some examples of security applications that run on Linux:

- ▶ **Symantec AntiVirus Scan Engine** provides advanced, high-performance virus scanning and repair services. It is designed specifically to protect traffic served through, or stored on, network infrastructure devices. The solution is easily integrated into server software applications and networking systems. Symantec AntiVirus Scan Engine detects and protects against viruses, worms, and trojan horses in all major file types, including mobile code and compressed file formats. For more information, see:

<http://enterprisesecurity.symantec.com/products/products.cfm?ProductID=173>

- ▶ **Norton AntiVirus for Lotus Domino** offers the most comprehensive, automatic protection against new and existing viruses. It keeps databases free from current and future viruses by automatically scanning and repairing file attachments in Notes mail and database documents. For more information, see:

<http://enterprisesecurity.symantec.com/products/products.cfm?ProductID=143&EID=0>

- ▶ **InterScan VirusWall** provides high-performance, comprehensive Internet gateway protection against viruses and malicious users of code. The optional eManager plug-in offers administrators additional tools to block spam, filter content, and schedule e-mail. For more information, see:

<http://www.trendmicro.com/en/products/gateway/isvw/evaluate/overview.htm>

- ▶ **InterScan WebManager** blocks unproductive Universal Resource Locators (URLs), monitors and manages Web usage, and scans Web traffic for viruses and other malicious code at the gateway. For more information, see:

<http://www.trendmicro.com/en/products/gateway/iswm/evaluate/overview.htm>

- ▶ **ServerProtect** software provides real-time antivirus scanning for Linux servers. It detects and removes viruses from files and compressed files in real time, before they reach the end user. Administrators can use a Web-based console to manage virus outbreaks, virus scanning, virus pattern file updates, and notifications. For more information, see:

<http://www.trendmicro.com/en/products/file-server/sp-linux/evaluate/overview.htm>

- ▶ **ScanMail for Lotus Notes** detects and removes viruses hidden in databases and e-mail in real time, before they spread to the desktop. It provides automatic virus signature file updates, centralized virus activity logs, and integration with the Lotus Notes administration client. For more information, see the following site:

<http://www.trendmicro.com/en/products/email/smln/evaluate/overview.htm>

- ▶ **FireWall-1** enables enterprises to define and enforce a single, comprehensive security policy that protects all network resources. Its

architecture delivers a highly scalable solution that integrates all aspects of network security. For more information, see:

<http://www.checkpoint.com/products/protect/firewall-1.html>

- ▶ **VPN-1 Pro** is the foundation of Check Point VPN-1 solutions, the most comprehensive set of products and technologies for remote access, intranet, and extranet VPNs. VPN-1 Pro software protects the privacy of information over the Internet while securing critical network resources against unauthorized access. For more information, see:

http://www.checkpoint.com/products/connect/vpn-1_pro.html

High-availability solutions

High-availability solutions are important to raise the availability of a system built on Linux. Such solutions include:

- ▶ **IBM Tivoli System Automation for Linux** supports the need for a high-availability solution that can also reduce complexity and costs. Failures, outages, and even slowdowns can result in direct business losses (the competition is just a mouse-click away) and may damage a company's credibility. The IBM autonomic computing initiative helps address these issues and more by using technology to make technology self-managing. IBM Tivoli System Automation for Linux provides self-healing capabilities for business applications by detecting failing IT components and then repairing or working around those failures.

For more information, see:

<ftp://ftp.software.ibm.com/software/tivoli/datasheets/ds-sys-auto-linux.pdf>

- ▶ **LifeKeeper for Linux** is a software application that ensures the continuous availability of applications by maintaining system uptime. LifeKeeper maintains the high availability of clustered Linux systems by monitoring system and application health.

For more information, see:

<http://www.steeleye.com/products/linux/>

- ▶ **VERITAS Cluster Server** eliminates both planned and unplanned downtime, facilitates server consolidation, and effectively manages a wide range of applications in heterogeneous environments. VERITAS Cluster Server supports up to 16 node clusters in SAN and traditional client/server environments.

For more information, see:

<http://www.veritas.com/products/category/ProductDetail.jhtml?productId=clusterserver>

Backup solutions

Although Linux controls backup and operating tape devices, additional backup solutions exist that make operations easy. They can also schedule backup and recovery events. Additional backup solutions include:

- ▶ **IBM Tivoli Storage Manager:** The latest version of IBM Tivoli Storage Manager server software is now available on Intel-based Linux systems. By extending IBM Tivoli Storage Manager's support for Linux from file systems and applications to include Tivoli Storage Manager servers, organizations using Linux environments can benefit from the same award-winning data protection as organizations using Tivoli Storage Manager servers on other platforms. IBM Storage Manager underscores IBM's commitment to provide robust, reliable software solutions for the mid-market through the large enterprise.

For more information, see:

<http://www.ibm.com/software/tivoli/solutions/linux/>

- ▶ **BrightStor ARCserve Backup for Linux** brings data protection technology to organizations deploying Linux. Beyond simple data protection, BrightStor ARCserve Backup uses multiple data verification methods to enable maximum data integrity and recovery capabilities.

For more information, see:

<http://www.ca.com/Solutions/Product.asp?ID=3370>

- ▶ **NetVault Enterprise Edition** is a comprehensive backup and restore solution for heterogeneous UNIX, Windows NT/2000, Linux, and Netware enterprise environments. It can rapidly add and configure new servers, devices, and clients, and can control them from a central location. As an organization grows, a modular architecture allows NetVault Enterprise Edition to combine with other BakBone plug-in modules, such as disaster recovery.

For more information, see:

http://www.bakbone.com/products/backup_and_restore/netvault_ee/

5.6.3 Customer scenario: A leading company in Web-based time- and expense- tracking solutions.

A company who is an IBM Business Partner and a leader in Web-based time- and expense-tracking solutions, provides time-, project-, and expense-tracking software to more than 10,000 organizations. Founded in 1996, the Austin, Texas-based company, employing approximately 50 people, already owns a 70% market share in its industry.

The company's flagship product, an online timesheet, that is integrated with IBM WebSphere Application Server and IBM DB2 Universal Database, enables

companies to efficiently track and manage project time and associated expenses over an intranet or the Internet.

A long-time DB2 user, the company based their timesheet on DB2 because of the data management system's consistently solid performance. The latest version of their timesheet leverages DB2 for Linux on the back end to maximize data access and management options.

In addition to Web-based applications, this company offers application hosting services to companies that don't want to manage the expense-reporting process in house. Although their timesheet supports a variety of databases, the company uses DB2 for Linux for all of its application hosting sites because of the robust data-management system of DB2 for Linux, which offers the flexibility, scalability, and pricing that support our business model.

The company hosts its timesheet application on nearly a dozen xSeries servers running Linux at its collocation facility. The company viewed Linux as a secure, robust, and cost-effective operating environment that would help it deliver reliable service. The xSeries servers also provide the high performance and reliability that the company needs. They also use an XML integration code that allows customers to customize the integration of multiple systems, including most well-known budget management, payroll automation, and accounting systems.

A leading professional services organization is a subscriber to the hosted Timesheet service. They have been a customer for over two years and rely on the time and tracking solutions company to accurately manage the timecard submission and approval process for a 900 employee department.

The process begins when an employee logs on to a secure Web page, served up by WebSphere Application Server, on the professional services organization's intranet. From the Web page, the employee completes the timecard and clicks a Submit button to route it to the backend server. From there, the time and tracking expenses solutions company's application collects the data, processes it, and routes it back to the professional services organization, using TCP/IP, with a copy stored in the DB2 database. Using this application, the professional services organization's managers access, review, and approve the timecards online. The data is then sent for payroll purposes, using an internal company application and an Independent Software Vendor's (ISV) HR module.

The company's timesheet has given the professional services company a reliable data access and management solution that saves the company time and money. The company has gained benefits from the online timesheet, which is generating an average saving of \$20,000 weekly or more than \$1 million annually.

By maximizing its operational efficiencies with Linux and DB2 and by managing the system implementation process, the Web-based time- and expense-tracking solutions passes substantial cost savings on to its customers. According to the company, most customers realize 100% payback within three to six months after installing the timesheet. They also eliminate the three to 10% error rate associated with manual timecard processing routines.

For more information about this success story, see the following site:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5F7T8L?OpenDocument&Site=db2software>

5.6.4 Customer scenario: LOSSRUNS Inc.

LOSSRUNS, Inc. (LRI) is reinventing how the property/casualty insurance industry distributes and accesses claim information. LRI disseminates information quickly and accurately to the people who need it by combining technological advances with industry expertise. LRI's patent-pending LOSSRUNS Report provides agents, carriers or underwriters with complete, correct and current claims information, which includes everything they need to properly evaluate risk and/or quote coverage.

LRI operates independently of any insurance company or related organization, supporting the independent agent system and the rights of both incumbent and non-incumbent agents.

The company's main goal was to have its application become a standardized solution to transmit historical claims data for the Property and Casualty insurance industry. The main problem was that potential customers wanted to know where and how it was going to store and distribute data in a secure and reliable fashion. LRI needed a strategic business partner to help facilitate its niche insurance software application and provide a solution that would help ensure data integrity on a stable, accessible platform.

LRI is using two IBM eServer xSeries 342 servers as the company's backup, application and database servers, as well as one xSeries 330 as the Web server. The xSeries 342 servers are running Red Hat Linux V7.2 and DB2 with multiple failover firewalls. The historical claims data used for building reports is being managed by DB2 and stored on an IBM 3583 Ultrium Scalable Tape Library. The xSeries 330 server is running WebSphere Application Server Version 4.0 Advanced Edition to develop the company's Web service interface with potential customers.

Property and Casualty (P&C) companies require historical claims reports to be generated for all customers when buying new or updating insurance policies. This was previously a cumbersome task and LOSSRUNS Report is making this

requirement less labor intensive while streamlining the process and making it less expensive.

At the IBM Atlanta hosting facility, IBM Global Services - e-business Hosting™ Services supports the configuration, offering its comprehensive suite of managed hosting services with enhanced features for security, backup, reporting, administration and monitoring. The agreement is a standard three-year hosting contract, for which LRI pays a monthly service fee.

The leading-edge technology, along with the IBM commitment to security and availability, enables LOSSRUNS to deliver a service its customers can rely on to be safe, secure and reliable. By using multiple levels of security technology, IBM provides peace of mind to insurance carriers, underwriters, and agents across the globe.

For more information about this success story, see the following site:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5LU26F?OpenDocument&Site=linuxatibm>

5.7 Infrastructure solutions

Infrastructure solutions continues to offer the largest saving opportunity as customers move to Linux. The typical solutions in this area are Web servers, mail servers, file/print servers, and firewall servers. Most of the Linux distributions for the xSeries server contain such open-source applications as Apache, Samba, IMAP, OpenLDAP, Perl, SendMail, Squid, and PostgreSQL.

Samba

Samba is an open source/free software suite. It provides file and print services to Server Message Block (SMB) and Common Internet File System (CIFS) clients. It allows end users to access and use files, printers, and other commonly shared resources on a company's intranet or on the Internet.

Samba is often referred to as a Network File System (NFS) and can be installed on a variety of operating system platforms, including Linux, most common UNIX platforms, OpenVMS, and OS/2. Samba supports the following UNIX and UNIX-compatible machines using Windows NT or 2000-compatible file/print server software. It can be used freely under GNU General Public License (GPL).

Configuration and management: Samba's configurations file is written in the only text file (smb.conf). In the Linux-plus-Samba environment, Samba can be set up and managed remotely using Telnet or ssh.

There is an easy configuration tool for Samba named SWAT. This tool can set up Samba from a Web browser.

Look for Samba on the Web at:

<http://www.samba.org/>

Apache

Apache is the most famous Web server (HTTP daemon) in the world. According to research by netcraft, Apache has more than 60% of the market share for Web servers. You can learn more about this research on the Web at:

<http://www.netcraft.com/survey>

Apache has two versions: Apache 1.3 and Apache 2.0. Table 5-10 shows what changed from Apache 1.3 to Apache 2.0.

Table 5-10 Changes from Apache 1.3 to Apache 2.0

Change	Details
Adopting Multi Processing Module (MPM)	MPM can communalize other modules on all platforms
Adopting Apache Portable Runtime (ARP)	Development by ARP and C language enables it to perform the same program also on other platforms
Correspondence to a multithred	To request processing by the thread generation from the process generation
Support of a multi-protocol	Apache 2.0 can easily support HTTP, HTTPS, and a new protocol
Support of IPv6	At Apache 2.0, Web server can be made to correspond to IPv6 only by base installation
Adopting new build system	Installation of autoconf and libtool is necessary to install Apache 2.0

To learn more about Apache, see:

<http://www.apache.org/>

Sendmail

Sendmail is the most common Mail Transfer Agent (MTA) for UNIX. Sendmail is configured by modifying sendmail.cf (or sendmail.m4). Sendmail can be set up by detailed configuration, but it is complex for system administrators to modify sendmail.cf. The wrong configuration can cause a security hole.

You can find Sendmail on the Web at:

<http://www.sendmail.org/>

Because Sendmail is heavy and complex for historical reasons, and has had some vulnerabilities in the past, some users preferred at some time to use Qmail and Postfix, and kept to that choice.

Qmail

Qmail is a secure, reliable, efficient simple message transfer agent. It is designed for typical Internet-connected UNIX hosts. As of October 2001, Qmail is the second most common Simple Mail Transfer Protocol (SMTP) server on the Internet and has by far the fastest growth of any SMTP server. The features of Qmail include:

- ▶ **Secure:** Security isn't a goal, but an absolute requirement. Mail delivery is critical for users. It cannot be turned off, so it must be completely secure. Qmail divides its functions into components. The executive user of each function is defined separately so that, if one component is damaged, other components are not influenced.
- ▶ **Reliable:** Qmail's straight-paper-path philosophy guarantees that, after a message is accepted into the system, it is never lost. Qmail also optionally supports maildir, a new, super-reliable user mailbox format. Maildir, unlike mbox files, is not corrupted if the system crashes during delivery. Even better, a user can safely read their mail over the NFS and any number of NFS clients can deliver mail to them at the same time.
- ▶ **Efficient:** On a Pentium under BSD/OS, Qmail can easily sustain 200,000 local messages per day. These are separate messages injected and delivered to mailboxes in a real test. Although remote deliveries are inherently limited by the slowness of DNS and SMTP, Qmail overlaps 20 simultaneous deliveries by default.

You can find Qmail on the Web at:

<http://www.qmail.org/top.html>

Postfix

Postfix software was developed as Sendmail-compatible MTA. Postfix has many features such as higher compatibility with Sendmail, security, easy maintenance, quick processing, and so on. Recently Postfix was adopted as default MTA by some Linux distributions. Postfix has the following merits:

- ▶ **Sendmail compatibility:** Postfix is similar in function and use as Sendmail. This makes education for new users and teaching them to set up new software for Postfix unnecessary. Of course, a minimum knowledge of MTA is essential. Shifting Postfix from Sendmail can be carried out without stopping its services.

- ▶ **Easy configurations:** It is too difficult for an average user to modify `sendmail.cf` manually. Usually a configuration tool, such as CF, is used. But the tools do not always follow the latest MTA, and there have been some incidents because of the bugs in the tools. The configuration file of Postfix can be modified easily with normal MTA skills.
- ▶ **Stability:** Postfix can run with parallel processing, so it is stable for business use.

You can learn more about Postfix on the Web at:

<http://www.postfix.org/>

Netfilter

Linux uses iptable or ipchains as a Netfilter of open source software. Generally, the Linux distribution of kernel-2.2x includes ipchains and that of kernel-2.4x include iptables. Netfilter analyzes the inside of TCP and IP header and filters the packet by referring to a filtering rule. When it filters, it checks:

- ▶ Source IP and destination IP
- ▶ Source port and destination port
- ▶ Protocol
- ▶ Interfaces
- ▶ ICMP type and code
- ▶ Checking SYN packet
- ▶ Checking fragment packet

In addition, iptable analyzes the following points:

- ▶ Stateful inspection
- ▶ TCP flags
- ▶ MAC address
- ▶ Limit of matching packet
- ▶ Owner of packet producer

To learn more about Netfilter, see:

<http://www.netfilter.org/>

Squid

Squid is the most common proxy server in Linux. Instead of a client, a Squid server accesses the server on the Internet. The Squid server relays the data and passes it to a client. Some advantages of using Squid are:

- ▶ You no longer need a client to access the Internet directly, and you can control access to the Internet.
- ▶ You can decrease network traffic by having Squid cache data.

- Since all traffic passes through Squid, you can match and investigate data.

You can find out more about Squid on the Web at:

<http://www.squid-cache.org/>

5.7.1 Customer scenario: Butterfly.net

Butterfly.net is a development studio and technology infrastructure provider for massively-multiplayer online PC, console and mobile games. With its unique server grid technology, Butterfly.net developed the best solution to the challenges of scalability, reliability and performance for persistent-state 3D worlds that must support millions of gamers.

Butterfly.net deployed the first commercial server grid for Massively-Multiplayer Games (MMGs), the most exciting and fastest-growing sector of the powerful video game market, in which millions of gamers worldwide face challenges together in real-time 3D worlds. Its MMG platform, dubbed the Butterfly Grid, is available as a fully-managed service, as a packaged software license, or as a complete hardware/software solution. Butterfly.net's Grid is the first system to allow MMG game worlds to scale in size and support an unlimited number of players within the same game. Grid does this by allocating communications and computing resources to the most populated areas and most popular games.

Butterfly.net was looking to launch its Butterfly Grid, an end-to-end solution for online video game developers, publishers and service providers, and the first commercial computing grid for the video game industry. Previously, Butterfly.net was experiencing slow response times due to too many players joining or gathering in one region. The company wanted to conquer this challenge, and saw an opportunity to leverage innovative grid computing technologies by using IBM Hosting Services.

Butterfly.net and IBM Global Services - Worldwide Strategic Outsourcing's (SO) e-business Hosting teamed to develop a video gaming grid powered by IBM e-business infrastructure housed at IBM e-business Hosting Centers. Linux-based IBM eServer xSeries 330 servers are used to run the game, gateway and daemon servers, and the xSeries 232 (also Linux-based) is used to run the DB2 Software. With the help of IBM, Butterfly.net developed the Information Grid in just two years and launched it rapidly, with limited risk and a manageable investment.

With the help of IBM, Butterfly.net has implemented a new infrastructure hosted by SO that has rapid time-to-market, lower IT costs and greater competitive advantage. Today, virtually everyone from developers to service providers to players benefits from the Butterfly Grid, which provides participants with reliable, high-performance and cost-effective capacity on demand. Game developers who

use the Butterfly.net software development kit, as well as publishers and service providers, no longer have to build a dedicated infrastructure to support a new game, reducing their risk and increasing their profit potential. For game-players themselves, the benefit is a better, faster, “always on” game experience.

For more information about this success story, see the following site:

<http://www.ibm.com/software/success/cssdb.nsf/CS/LE0D-5MQTPP?OpenDocument&Site=default>

5.8 xSeries architected solutions

IBM provides a comprehensive solution for modular computing through the blade architecture. The blade architecture is built from blade servers and BladeCenters.

5.8.1 Blade server

The rapid spread of the Internet over several years has pushed the server system that supports it. The scale of e-business expanded quickly. The typical system complex is the high-density mounting of servers, built from the 1U rack-mounted server. In 2001, to meet the demand for super-high-density mounting, a new server appeared, the blade server.

The term *blade* means a flat section, like a knife blade. The blade server is a long, slender circuit board with condensed server functions. Equipped with six to 20 blades in a case called the *enclosure* that is 2U-7U size, it achieves a density greater than the 1U server.

BladeCenter

In October 2002, IBM announced the Intel Xeon processor two-way SMP called the blade server *BladeCenter HS20* and delivered the *BladeCenter* enclosure, which was 7U size. BladeCenter and BladeCenter HS20 are designed to save space in a high-density rack environment. BladeCenter can store 14 BladeCenter HS20 blades with a maximum of six BladeCenters in the 42U rack. BladeCenter HS20 supports Red Hat Linux and SuSE Linux.



Figure 5-4 BladeCenter

Characteristics of the BladeCenter marketplace include:

- ▶ **Server consolidation:** BladeCenter is a great solution for a centralization-type server consolidation.
- ▶ **HPC/grid computing:** BladeCenter is extremely good for HPC environments with the highly efficient blade Intel Xeon processors running the following workloads:
 - Linux cluster
 - Genome
 - Engineering Design Automation (EDA)
 - Grid
- ▶ **Application server:** Since BladeCenter has the same performance and function as a general-purpose server (for example, a SAN environment), the BladeCenter becomes the optimum application server and can provide considerable space savings in the following environments:
 - File server
 - Notes server
 - WebSphere Application Server

For more information, see the following Web sites:

http://www.pc.ibm.com/us/eserver/xseries/bladecenter_family.html

<http://www.ibm.com/servers/eserver/blades/>

5.8.2 Customer scenario: Wolferman's

Wolferman's has been a proud purveyor of specialty foods for more than 100 years, beginning with the opening of the first Wolferman's grocery store in Kansas City, Missouri in 1888. Today, Wolferman's also operates a direct-mail specialty catalog company, which typically generates 80% of its annual sales between mid-October and the end of December each year.

The Wolferman's Internet site receives the majority of its traffic during the holiday season each year when shoppers nationwide flock to the Internet to purchase gift

baskets. And while Wolferman's may be renowned for its delicious recipes, the online gift basket business is highly competitive, so Web site performance is critical to customer satisfaction. Wolferman's was experiencing significant growth in its e-commerce sales, but its previous Web site infrastructure could not scale to keep up with the growing business. As a result, its online catalog solution was periodically unable to support customer demands.

Wolferman's customers could also order the company's products through a toll-free call center, but an average order placed through the Wolferman's 800 number costs the company 80% more than orders placed over its Web site. The company calculated that it could save hundreds of thousands of dollars per year by encouraging customers to make purchases over the Web rather than over the phone. Consequently, based on this combination of factors, Wolferman's wanted to significantly improve its e-commerce capabilities in an effort to improve customer satisfaction and loyalty while optimizing costs and competitiveness.

Previously, Wolferman's e-commerce solution consisted of third-party applications running on an IBM AS/400 server. The solution architecture made it difficult for the company to support growth and processing speeds, and reliability was becoming a problem. Moving forward, Wolferman's wanted to implement a more robust, scalable, and manageable e-commerce solution that would enable it to grow its Web business in line with customer demands.

Wolferman's selected an "out-of-the-box" Web site solution from an IBM Business Partner, eOne Group, which is built on a foundation of IBM hardware and software. The new e-commerce site includes the following key components:

- ▶ IBM eServer Integrated Platform for e-business
- ▶ IBM WebSphere Application Server - Enterprise Version 5.0
- ▶ IBM DB2 Universal Database for Linux
- ▶ IBM DB2/400 for iSeries
- ▶ IBM eServer xSeries and iSeries servers
- ▶ SuSE Linux
- ▶ eOneCommerce
- ▶ eOneConfigurator

By selecting eOne Group's eOneCommerce running on an IBM infrastructure, Wolferman's was able to quickly implement and deploy an affordable, feature rich, and fully integrated Web site. eOneCommerce is an enterprise Java, SQL, and Extensible Markup Language (XML)-driven e-commerce and portal application. It provides companies with virtually all of the standard functionality needed for an e-commerce site, including unlimited browsing categories, membership, searching, order status and history, promotions, item links and relationships, favorites, personalization, and flexible pricing.

Because eOneCommerce is a 100% Java application, it requires an application server. In this case, WebSphere Application Server provides eOneCommerce support for all of the consumers that shop at <http://www.Wolfermans.com> including memory management of product information, images, and descriptions, as well as session management to track items that are in shoppers' carts. The WebSphere technology also supports category browsing, product searches, and gift basket customization options. The membership functionality within eOneCommerce also leverages WebSphere to provide Wolferman's customers with the option to save multiple ship-to addresses in their profile along with their history and favorites list.

Two IBM @server xSeries 330 servers, which are running DB2 and WebSphere and are powered by the Linux operating system, are used as the Web-serving platform. Wolferman's selected the x330 servers for their scalability, speed, and performance capabilities given that Web application server performance and reliability were business-critical for the new e-commerce solution. The x330 servers are also integrated with Wolferman's backend iSeries server, which is used to manage and deliver critical transaction data.

One employee now maintains the Wolferman's Web site, and it can support up to 6,000 concurrent users during the holiday season.

The combination xSeries, WebSphere, and eOneCommerce solution provided Wolferman's with an affordable, feature-rich, fully integrated Web site that is powerful and scalable enough to support dramatic seasonal sales spikes. With fewer infrastructure requirements, the new solution also has a lower total cost of ownership (TCO) than the previous solution. For example, Wolferman's estimates that the new e-commerce site will save the company more than US \$200 thousand per year in hardware, software, and management costs.

Since implementing the solution, Wolferman's has also been pleased with the performance and stability of the Linux-based x330 servers. Its Web pages now load much faster - even when the site is experiencing peak holiday traffic - and site downtime has been virtually eliminated. As a result, user customer satisfaction has increased, which translates into fewer lost sales and more repeat customers.

For more information about this success story, see the following site:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5MD35R?OpenDocument&Site=default>

5.9 xSeries services solutions

Specific types of service vary by geographical area, but it generally includes implementing Linux solutions for new customers and providing support for existing customers. The following sections pertain to service in the United States. For more information, see:

<http://www.ibm.com/services/e-business/linux.html>

5.9.1 Implementing Linux

The implementation of Linux service ranges from installation support to the construction of clusters. Installing and configuring Linux are important business opportunities, as are setting up Apache and Samba. However, since many companies offer those basic services, the price of such services (and the profit from providing them) is decreasing. High-added-value services such as cluster construction and server consolidation will be an even more important business opportunity for IBM in the future.

The following sections explain the typical types of service available for Linux on xSeries.

Linux Solution for e-business

The Linux Solution for e-business provides rapid deployment of a customer's Linux-based e-business environment. This solution is tailored to fit the specific needs of the customer by allowing the customer to select from a wide variety of middleware and applications, which can be installed and configured on the base of an IBM-supported Linux operating system.

The Linux Solution for e-business service capability consists of two phases. Each phase consists of a set of selectable elements.

- ▶ **Phase 1, Proof of concept:** This phase involves installing, customizing, and demonstrating both base and selected additional technologies in an environment isolated from the customer organization's existing production systems.
- ▶ **Phase 2, Going to production:** This phase involves the following elements:
 - Assisting the customer organization in connecting the newly installed e-business server to its existing systems
 - Establishing the Linux Solution for e-business server connectivity to a public or private TCP/IP network
 - Connecting one of the customer organization's existing database or transaction manager subsystems to the corresponding subsystem on the

Linux Solution for e-business server, or implementing collaboration and Internet servers.

For more information, see:

http://www.ibm.com/services/e-business/linux_3.html

High-performance clusters services

For customers who design, implement, and support Linux clusters for Internet and enterprise-wide, mission-critical applications, IBM is there to help.

To learn about these services, see:

http://www.ibm.com/services/e-business/linux_6.html

Server consolidation service

The driving force behind server consolidation is the need to reduce TCO. Many customers have a mishmash of hardware and operating systems that are expensive to operate. Also, more modern applications such as enterprise-resource-planning software further complicate matters. Using an architecture-independent approach, IBM helps customers reduce the TCO.

For more information, see:

http://www.ibm.com/services/e-business/linux_7.html

5.9.2 Support

Depending on customer need, IBM offers Internet and voice support, ranging from answering usage questions to identifying problems. This support is available 24 hours a day, 7 days a week. IBM Global Services also provides consulting, planning, and implementation services for Linux. IBM consultants can help you evaluate whether Linux is appropriate for your particular environment.

For more information, see:

http://www.ibm.com/services/e-business/linux_8.html



pSeries and Linux

Linux on pSeries is a key part of the IBM Linux strategy. During the past few years, IBM has been developing the ability to integrate Linux capabilities into the pSeries server. pSeries servers can run Linux alone, or can run both IBM AIX and Linux in logical partitions (LPAR) of the same server.

IBM as a corporation has embraced Linux as if it were one of its own internally developed operating systems. Linux receives full support from all of IBM's software development organizations including Lotus, Tivoli, the DB2 products, and the WebSphere products, and all of these software products are available for Linux and are fully supported. IBM's Server division has enabled Linux on all IBM @server platforms from the Intel-based xSeries to the PowerPC-based iSeries and pSeries, to the zSeries mainframe. And IBM has a Linux Technology Center staffed with more than 250 IBM programmers and engineers working full-time on various open source projects such as Linux. IBM's Services division offers training, consulting, and support to its customers. There is no doubt from any quarter that the entire IBM corporation world-wide is fully committed to Linux.

This chapter provides information about Linux enablement for the pSeries. It discusses the ways that you can integrate Linux applications into commercial pSeries environments.

6.1 How Linux runs on the pSeries server

The pSeries server is one of IBM's key solutions to fit a variety of customer requirements. pSeries servers offer leading edge performance, and can help you confidently manage your business's growth and risks.

To meet your requirements, pSeries products offer a full range of high-performance servers complete with highly functional, state-of-the-art software. IBM will provide more services for Linux for pSeries in response to growing interest from customers to use Linux as their operating system in this platform.

There are some key concepts worth reviewing before you see how Linux and AIX systems and applications run on the pSeries. For more information, see 1.3, "Linux: The operating system" on page 14.

For more information, see IBM's Linux for pSeries Web site:

<http://www.ibm.com/servers/eserver/pseries/linux/index.html>

6.1.1 Linux-capable pSeries model structure

The Linux-ready pSeries server has models that range from entry desk side models through to high-end enterprise models. Rack mounted models offer physical consolidation possibilities. Figure 6-1 shows all-Linux capable models, from the low-end desk side and desktop models, to mid-range rack mounted models, to the high-end servers. The trend has been that all new pSeries models are Linux-ready from the very beginning. See IBM's pSeries Web page for current information on pSeries models, technical specifications, white papers, and other information:

<http://www.ibm.com/eserver/pseries>



Figure 6-1 pSeries models

The following pSeries white papers are a must-read for anyone who wants further information about Linux on pSeries:

- ▶ *IBM @server Linux for pSeries: An overview for Customers*
http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html
- ▶ *Linux for IBM @server pSeries: Facts and Features*
http://www.ibm.com/servers/eserver/pseries/hardware/linux_facts.pdf
- ▶ *LPAR for Decision Makers*
http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/lpar_decision.pdf
- ▶ *Linux RAS for IBM pSeries: Reliability, Availability, Serviceability (RAS)*
https://techsupport.services.ibm.com/server/Linux_on_pSeries/images/Linux_RAS.pdf
- ▶ A collection of pSeries white papers:
http://www.ibm.com/servers/eserver/pseries/library/wp_systems.html

For more information about pSeries logical partitioning, refer to the IBM ITSO Redbooks Web site at:

<http://www.redbooks.ibm.com>

6.1.2 pSeries operating system scenarios

The pSeries offers outstanding server performance and a great server consolidation platform. These servers provide the ability to run AIX or Linux as single operating systems, or multiple LPARs running any combination of AIX and Linux. Figure 6-2 shows an overview of the internals of the AIX and Linux operating systems, and how applications can run on pSeries in three different scenarios.

Scenario 1 shows the traditional AIX-only-based pSeries environment. All AIX versions prior to Version 4.3.3 or 5L only run applications compiled for AIX. That means Linux applications cannot be run, unless they are explicitly ported to AIX.

Scenario 2 shows how AIX 4.3.3 and AIX 5L™ incorporate Linux in the pSeries by means of the *AIX Toolbox for Linux Applications*. The AIX Toolbox for Linux Applications provides the capability to recompile open source applications to run as AIX applications. In most cases, it does this without porting and using the same open source build environment as Linux.

Scenario 3 shows the fully Linux-enabled scheme. It has a Linux for pSeries kernel. Customers can integrate a broad spectrum of Linux applications into their pSeries servers. The customer can generate Linux code for any Linux application available for the pSeries.

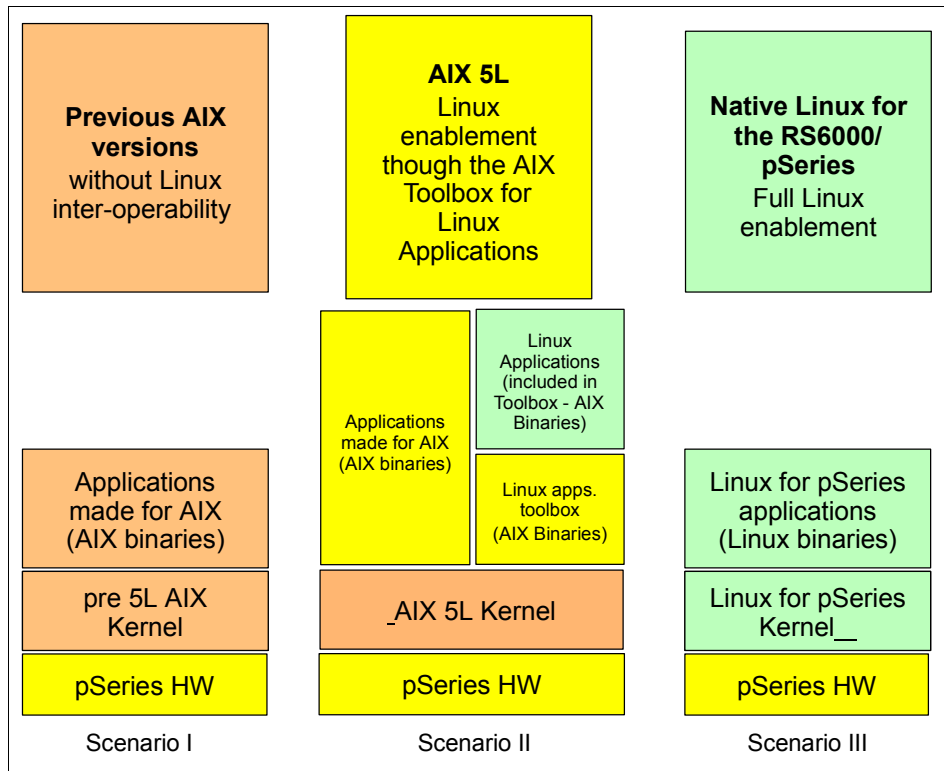


Figure 6-2 How UNIX operating systems and applications run on the pSeries

6.1.3 AIX 5L and Linux

You can run many Linux applications on AIX, a mission-critical, highly scalable operating system through Linux application interoperability to AIX 5L.

IBM has provided the AIX Toolbox for Linux applications, which is supplied as a supplementary CD with the AIX Version 5L operating system. This toolbox is also available for downloading. You will find a link to the download from this Web site:

<http://www-1.ibm.com/servers/aix/products/aixos/linux/index.html>

You can find more information about AIX by going to the main AIX Web page:

<http://www.ibm.com/AIX>

AIX Affinity with Linux uses an API approach to provide Linux application interoperability. This approach is *not* an environment nor an additional layer nor a

wrapper on or in which to run Linux applications. It is the integration of Linux compatible APIs and header files into AIX 5L.

Therefore, recompiled Linux applications become AIX applications, and have access to all the reliability, scalability, and availability of AIX. The result is tighter integration of the application to the operating system than can be achieved with an Application Binary Interface (ABI) approach.

AIX was developed using UNIX industry standards. As such, there is a high degree of compatibility at the API level between AIX and Linux. This degree of similarity allows many Linux applications to be recompiled and run on AIX 4.3.3 and AIX 5L using the AIX Toolbox for Linux Applications.

AIX 5L Version 5.2 Linux affinity provides an integrated Linux application environment within the base AIX 5L operating system with the objective to support “compile and go” enablement for Linux applications developed on Intel® Architecture systems. It enables faster and less costly deployment of multi-platform, integrated solutions across AIX 5L and Linux platforms. Through Cluster Systems Management (CSM), AIX 5L supports both Linux and AIX 5L heterogeneous distributed cluster systems management. AIX 5L also offers a Web-based system management client specifically for Linux. More information about AIX 5L Linux affinity can be found in the white paper *IBM @server pSeries AIX 5L v5.2, the On Demand UNIX Operating System for POWER Platforms*, at this Web page:

http://www-1.ibm.com/servers/aix/whitepapers/aix_ondemand.pdf

IBM provides the AIX® Toolbox for Linux Applications, a collection of Open Source and GNU software commonly found with Linux distributions. Over 360 separate development tools, server applications, and Linux utilities are available in both open source code and binary Red Hat Package Manager (RPM) format for AIX 5L. The AIX Toolbox for Linux Applications is available at no charge on CD-ROM or from the following IBM Web site:

<http://www.ibm.com/servers/aix/products/aixos/linux/>

You can learn more about AIX Affinity in the white paper at:

http://www.ibm.com/servers/aix/products/aixos/linux/affinity_linux.pdf

The uses of AIX Toolbox for Linux Applications include:

- ▶ Build and package Linux applications for use on AIX
- ▶ Run Gnome and KDE desktops
- ▶ Run other popular software commonly found in Linux distributions

- Manage open source software using the popular RPM package management system
- Develop new applications for AIX using GNU and Linux application development tools

Figure 6-3 shows the AIX toolbox structure running over the AIX kernel.

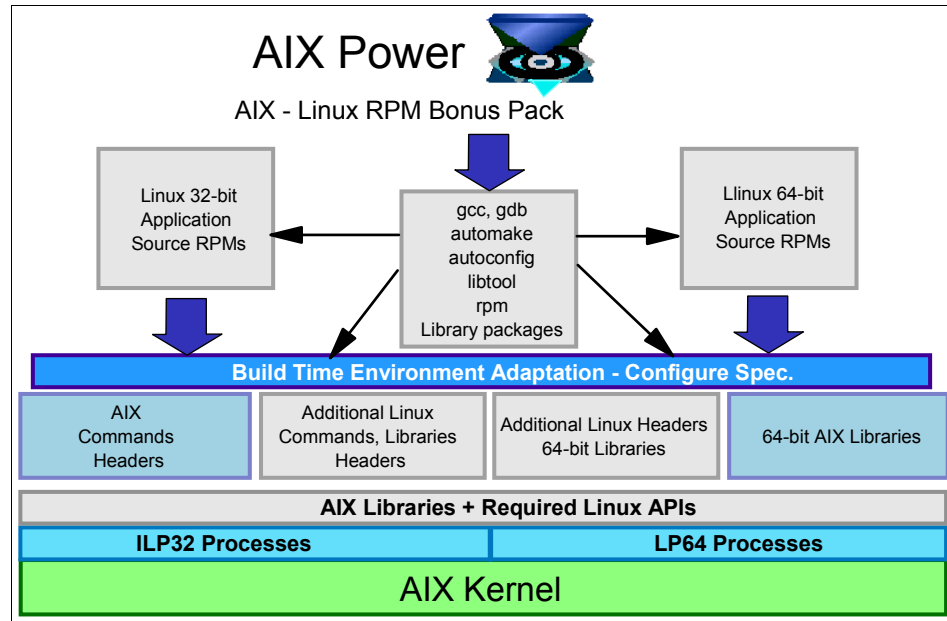


Figure 6-3 The AIX Toolbox for Linux Applications

The AIX Toolbox for Linux Applications contains a wide variety of software, including:

- **Application development:** gcc, g++, gdb, rpm, cvs, automake, autoconf, libtool, bison, flex, gettext
- **Desktop environments:** GNOME and KDE
- **GNU base utilities:** gawk, m4, indent, sed, tar, diffutils, fileutils, findutils, textutils, grep, sh-utils
- **Programming languages:** guile, python, tcl/tk, rep-gtk
- **System utilities:** emacs, vim, bzip2, gzip, git, elm, ncftp, rsync, wget, lsof, less, samba, zip, unzip, zoo
- **Graphics applications:** ImageMagick, transfig, xfig, xpdf, ghostscript, gv, mpage, the Gimp

- ▶ **Libraries:** ncurses, readline, libtiff, libpng, libjpeg, slang, fnlib, db, gtk+, qt
- ▶ **System shells:** bash2, tcsh, zsh
- ▶ **Window managers:** Enlightenment, sawfish

For the latest information about AIX-5L, visit the following Web site:

<http://www.ibm.com/servers/aix/os/52desc.html>

AIX 5L also introduces many new features. You can learn about them in the IBM Redbook *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765.

Exploring the details of running Linux applications on AIX as well as other built in functionality is beyond the scope of this book. We recommend that you consult the redbook *Linux Applications on pSeries*, SG24-6033.

You can find more information about the AIX Toolbox for Linux Applications on the Web at:

<http://www.ibm.com/servers/aix/products/aixos/linux/index.html>

AIX 5L is intended to meet the need for providing interoperability between Linux applications (as well as other Linux machines) and the pSeries that run AIX. Bringing Linux applications to the pSeries with AIX 5L makes interoperability a reality. In addition, you can retain the considerable benefits and functionality derived from AIX.

6.1.4 Options to run Linux on the pSeries server

Customers have many choices with the pSeries servers. They can run Linux applications under AIX, or run Linux as an operating system per server (similar to unpartitioned Intel servers). Or Linux can run as a partition among other Linux or AIX partitions. There are many variations possible. This gives you a great deal of flexibility to manage and possibly optimize your hardware resources.

This section covers both non-partitioned and partitioned Linux servers.

Linux on the pSeries

We assume that you have selected a Linux application that does not suit the AIX-5L approach. Plus you want to run the application on a separate server. This can be indicative of the real environment where a dedicated server is installed to run specific a application.

Linux applications can run on pSeries LPAR hardware by simply installing the Linux operating system.

While the AIX Toolbox for Linux Applications offers a way to port Linux applications to pSeries servers running AIX, there may be cases when it is preferred to run them under a Linux operating system instead of AIX.

The process is simple with the following steps:

1. Select the Linux application that is supported by one of the Linux Distribution Partners (LDPs). Linux applications are frequently available as a demo or trial download for testing.
2. Size and select any pSeries models that supports Linux and which meets your needs from one of the Linux Express or other pSeries models.
3. Then you order the hardware, Linux distribution, Linux application, and necessary support contracts.

These Linux Express models are not shipped with the Linux operating system pre-loaded. The Linux Express models have an optional feature that includes a particular Linux distribution with your order. When you select this optional feature, IBM will send the order to its Linux Distribution Partner (LDP) for fulfillment. You are not required to order AIX for this Linux Express solution.

This approach also applies to any other Linux applications that a pSeries customer may run, which are not available with the AIX Toolbox for Linux Applications.

Linux LPAR on a supported pSeries server

The LPAR capabilities of some pSeries servers make it possible to run one or more instances of Linux along with AIX. This provides a convenient way to begin developing and deploying Linux operating system-ready applications as desired, while retaining the enterprise-ready capabilities of AIX for mission-critical or highly-scalable workloads. LPAR also allows large pSeries servers to be partitioned to run Linux only workloads.

First, let us look at some terminology:

- ▶ **LPAR:** This is the shortened version of the familiar term *logical partitions*. It is widely used in zSeries, iSeries, and pSeries servers to indicate that a server can run multiple separate operating system instances.
- ▶ **Static LPAR:** This term applies to a pSeries server or partition. It indicates that the server does not support the rearrangement of resources without powering down. The same restriction applied to early releases of iSeries LPAR. All pSeries Linux partitions are static.
- ▶ **Dynamic LPAR (DLPAR):** This term indicates that a partition accepts certain rearrangements of resources. Partitions running AIX version 5.2 applications can have dynamic movement of processors and memory between partitions.

- ▶ **Hypervisor:** This is the underlying control mechanism that sits between the physical hardware and the operating systems. The hypervisor owns all system resources, and creates partitions by allocating these resources and sharing them.
- ▶ **Hardware Management Console (HMC):** This is the computer (an Intel-based PC running Linux) that provides the low-level hardware interface to the pSeries server through a serial link. It is through the HMC that the administrator creates and modified LPARs, and other low-level hardware control functions. Higher level control functions can be accessed through an Ethernet link once the LPARs have been booted with an operating system.

AIX is not required to partition the machine, and it is not required to install and run Linux. Currently, AIX provides greater features for dynamic partition support, full 32-way SMP support, hardware analysis, intermittent fault detection, and error logging.

For more information, visit the pSeries Web site at:

<http://www-1.ibm.com/servers/eserver/pseries/linux/>

AIX and Linux can run in separate partitions at the same time on an LPAR-enabled pSeries server in any combination (that is, zero or more Linux partitions along with zero or more AIX partitions). This allows you to consolidate workloads from several separate servers onto a single system.

Since partitioning in an LPAR-supported pSeries is controlled by the firmware-based hypervisor and the Hardware Management Console (HMC), AIX is not required to run Linux.

Figure 6-4 shows a typical logically partitioned pSeries server. It shows the major components, the HMC linked both to the partitions, and to the hypervisor. This includes the internal components that the HMC can control, processors, memory, and I/O adapters.

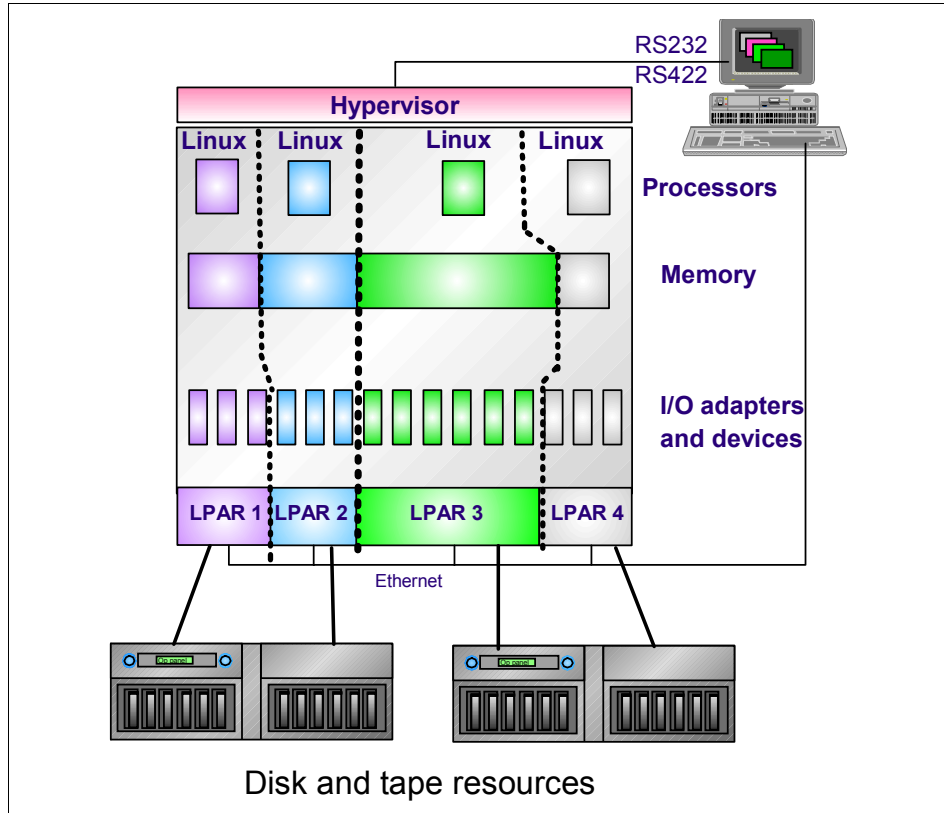


Figure 6-4 Linux for the pSeries server on the LPAR

The logical partitioning functions allow the rearrangement of resources to and from a partition, as well as to move resources between two partitions, without operating system reboot. The new Dynamic LPAR function in pSeries can dynamically move resources for AIX 5L Version 5.2 partitions. It is currently not supported by Linux or earlier versions of AIX. However, Linux partitions and dynamic LPAR can run on the same server.

For more information about LPAR-supported pSeries servers, see:

<http://www.ibm.com/servers/eserver/pseries/lpar/resources.html>

In the example in Figure 6-4, any partitions that are shown as running Linux can also be shown as a partition running AIX.

6.1.5 Linux and its relationship with AIX

Linux can relate with AIX in two ways. There may be system environments in which both types of relationships are present at the same time:

- ▶ The ability to run Linux applications ported to AIX 5L
- ▶ The ability to have Linux-based pSeries front end systems interacting with mission-critical AIX-based pSeries servers

Linux has already been helping AIX, as AIX 5L has gained more momentum due to the AIX Toolbox for Linux Applications, and the increasing built-in Linux affinity of AIX. By allowing Linux applications to be easily ported to AIX with the AIX Toolbox for Linux Applications package, IBM is expanding the available application portfolio. In addition, AIX can now take advantage of the growing Linux skill base.

Furthermore, Linux and AIX are complementary operating systems. AIX is the strategic, proven, mission-critical operating system for the pSeries. Linux is a highly portable operating system that supports all IBM platforms. IBM expects to see many installations running Linux (on xSeries or pSeries hardware) as the front end to mission-critical AIX systems running DB2 and other enterprise applications.

As you can see, AIX is the strategic operating system for the pSeries server. It contains such enterprise features as scalability, advanced performance characteristics, and reliability, availability, and serviceability (RAS) features. The high-end enterprise versions of Linux are rapidly closing the gap with the release of new function in new versions of these Linux distributions for pSeries servers.

6.2 Linux for pSeries hardware enablement

Both 32-bit and 64-bit versions of Linux for pSeries are being provided to optimize customer choices and exploit pSeries hardware capabilities. The 64-bit POWER3™ and POWER4™ systems have a 32-bit or 64-bit kernel, and support a 32-bit or 64-bit application environment (depending on the model and Linux distribution). Current Linux development efforts are focused on the 64-bit products.

I/O device and adapter support and availability

Many adapters and devices can be attached to pSeries systems. While some of the devices (for example, PCI adapters) have Linux drivers for Intel and AIX drivers, they cannot be used directly in Linux for pSeries.

However, Linux for pSeries has base device support. Be sure to review the latest version of the *Linux for pSeries Facts and Features* whitepaper on the Web at:

http://www.ibm.com/servers/eserver/pseries/hardware/linux_facts.pdf

Important: Some of the hardware devices listed on any LDP site may not be fully enabled or certified. We encourage you to look at each LDP's Linux hardware support database, and compare it with the IBM Support matrices.

Should you encounter problems with unsupported graphics adapters, unplug the card and use a serial console to install Linux. We encourage you to periodically look for updated information about the hardware platforms in which Linux for PPC is intended to be enabled and has limited or full support. See:

<http://www.ibm.com/servers/eserver/pseries/linux>

6.2.1 Linux for pSeries scalability

There is an excellent publication, *Towards Linux 2.6: A look into the workings of the next new kernel* by Anand K Santhanam of IBM India. This publication is available on the developerWorks Web site at:

<http://www.ibm.com/developerworks/linux/library/l-inside.html>

This article discusses in detail the scalability of the Linux 2.6 kernel, and how the new kernel differs in scalability from the 2.4 kernel. With regard to scalability on SMP systems, the article describes the how the scheduler works in the 2.4 kernel and the improvements made in the 2.6 kernel. The new 2.6 kernel has redesigned its SMP scheduling algorithms and is expected to perform well in 8-way to 16-way and larger SMP environments. Some of these 2.6 kernel features have been back-ported to the 2.4 kernel. This publication contains many links that probe more deeply into detailed areas relating to the main article.

Other than Santhanam's article cited above, there is little or no specific published information on the scalability of Linux on the pSeries or any other platform -that is, the N-way SMP capabilities of the kernel on pSeries servers.

After an exhaustive search of the available literature and Web sites both within IBM and with the Web sites of the Linux distributions, we could not find any specific published information on Linux scalability. Most of the claims of scalability of Linux seems to be associated with Linux clusters of one to two-way machines, rather than in a Linux implementation on a single larger (8 to 16-way) SMP server.

6.2.2 Linux for pSeries's Reliability, Availability, and Serviceability

A key attribute of the pSeries is its critical Reliability, Availability, and Serviceability (RAS) features. RAS is a concept used not only by IBM, considering that RAS features are the same. However, implementation may vary depending on the platform and environment on which you are working. Linux for pSeries is rapidly developing RAS capabilities. However, many pSeries RAS features are only fully realized when running AIX. For more information on RAS features on pSeries running Linux, see the white paper at this URL:

https://techsupport.services.ibm.com/server/Linux_on_pSeries/images/Linux_RAS.pdf

The following pSeries RAS features are supported when running Linux:

- ▶ Automatic First Failure Data Capture and diagnostic fault isolation capabilities
- ▶ Self-healing internal POWER4 processor array redundancy
- ▶ Scrubbing and redundant bit-steering for self-healing in main storage
- ▶ Chipkill™ and ECC correction in main storage
- ▶ Memory error correction extensions
- ▶ Fault tolerance with N+1 redundancy of power and cooling, dual line cords, and concurrent maintenance for power and cooling
- ▶ Predictive failure analysis on processors, caches, and memory
- ▶ Journaled file system (several available under Linux)
- ▶ PCI Extended Error detection
- ▶ PCI bus parity error recovery
- ▶ Slot freeze detection
- ▶ Unrecoverable error handling
- ▶ Fault avoidance through highly reliable component selection, component minimization and error mitigating technology internal to chips
- ▶ Concurrent run-time diagnostics for power and cooling
- ▶ Error logging and reporting
- ▶ Service processor (available now) and Service Agent (available 2H04)
- ▶ Service Focal Point on the Hardware Management Console (HMC)
- ▶ Boot-time processor and memory deallocation (Persistent Processor Deallocation)

Some of the pSeries RAS features that are not currently supported in Linux include:

- ▶ Hot-swapping of disk drives (this can be done through command line procedures to some degree)
- ▶ Hot-plug PCI and memory
- ▶ Disk mirroring
- ▶ Dynamic Processor Deallocation
- ▶ Dynamic logical partitioning (static LPAR is supported)
- ▶ Error log analysis

- ▶ Remote support
- ▶ High Availability Cluster Multiprocessing (HACMP™). Alternative third party and Open Source solutions may be available.

6.2.3 Other sources of information

Through its Linux Technology Center, IBM is working with the open source community on a variety of projects to enhance the value of Linux for customers. The Linux Technology Center has over 200 people devoted to developing and improving open source.

You can find the Linux Technology Center on the Web at:

<http://www.ibm.com/developerworks/oss/linux/>

IBM also participates in several industry-led efforts, such as:

- ▶ Linux Standard Base
<http://www.linuxbase.org>
- ▶ Free Standards Group
<http://www.freestandards.org>
- ▶ Open Source Development Network
<http://www.osdn.com>
- ▶ Open Source Development Lab
<http://www.osdlab.org>

PartnerWorld® Linux Lens

Sharpen your focus through the Linux Lens. The Linux Lens is IBM's most comprehensive resource for IBM Business Partners interested in Linux. Providing up-to-date news, tools and resources, the Linux Lens is designed to be your one-stop shop for IBM Linux Business Partners information. This Web site is a great source of information about conferences, programs, and events that are important for IBM customers, business partners, IBMers, and other interested parties:

http://www-1.ibm.com/linux/va_12.shtml

Here you will find information about major conferences such as the LinuxWorld Conference & Expo, PartnerWorld conference, and links to various other Linux related resources, distributions, white papers, and so on.

Education and conferences

In September 2004 IBM is hosting a conference called the *IBM Eserver pSeries and Linux Technical University*. There are links to other IBM conferences around the world. The Web site for this conference is:

<http://www-3.ibm.com/services/learning/conf/us/pseries/preconf.html>

This next URL takes you to one of the presentations from a previous pSeries and Linux Technical University conference:

[http://www.ibm.com/partnerworld/pwhome.nsf/vAssetsLookup/LunchLearnLinuxforpSeries2002_08_27.pdf/\\$File/LunchLearnLinuxforpSeries2002_08_27.pdf](http://www.ibm.com/partnerworld/pwhome.nsf/vAssetsLookup/LunchLearnLinuxforpSeries2002_08_27.pdf/$File/LunchLearnLinuxforpSeries2002_08_27.pdf)

Linux Standards Base (LSB)

Important advances have been made in the development of Linux code that is binary-compatible across different hardware platforms. The Linux Standards Base (LSB) represents a development standard and certification process to generate and run operating system and application binary code successfully on different platforms. IBM's Linux Technology Center is an active participant in the LSB technology. The Linux Technology Center Web site describes the Linux Standards Base, lists the certified distributions (seven distributions are certified as at this writing), and has links to the detailed descriptions of the exact standards and specifications that a piece of developed software must implement in order to achieve LSB certification. You can find the LSB page of IBM's Linux Technology Center Web site at this URL:

<http://ltc.linux.ibm.com/lsb/index.html>

The ultimate goal is for any Linux application to run on any Linux distribution utilizing the LSB's binary compatibility methodology. Testing and maintenance should be dramatically decreased while increasing the variety of Linux systems to run on. The IBM Linux TEchnology Center is available to assist IBM's Software Group with LSB compliance. IBM's LSB team can help SWG developers understand their current level of compliance, what it takes to achieve full compliance, and the actual process of becoming certified.

Linux kernel

The source code for the Linux kernel can be found at the Web site:

<http://www.kernel.org/>

This site is the authoritative repository for the source code for the Linux kernel for all platforms, including PowerPC 32-bit and 64-bit kernels. A first port of the Linux kernel to the PowerPC processor, which is the core of a pSeries system, was developed by the open source community a few years ago. IBM has since

taken an active, leadership role in the development of Linux for the PowerPC. The terms “PowerPC” and “PowerPC processor” are used in this book to refer to all IBM processors that are based on the PowerPC architecture (such as PowerPC 601, PowerPC 603™, PowerPC 604™, PowerPC 604e, RS64 II, RS64 III, RS64 IV, Power2, Power3, Power4).

Currently, the kernel.org Web site lists two main projects related to the Linux kernel on the PowerPC processor architecture: 32-bit and 64-bit. The *PowerPC* link takes you to the URL:

► <http://penguinppc.org/>

This project provides a 32-bit Linux kernel for PowerPC based systems. It can be used on a large variety of pSeries, iSeries, RS/6000 machines, Power Macintoshes, PowerBooks, iMacs and also the Motorola 8xx and IBM 4xx embedded processors. You can refer to these Web sites for a listing of supported systems:

<http://penguinppc.org/projects/hw/>

and

<http://oss.software.ibm.com/developer/opensource/linux/projects/ppc/>

The *PowerPC64* link in the kernel.org Web site takes you to the URL:

► <http://penguinppc64.org>

This project provides a 64-bit kernel for 64-bit PowerPC based hardware, including Power3 and Power4 based IBM pSeries and iSeries systems, and the BladeCenter JS20 server.

Please see

<http://oss.software.ibm.com/developer/opensource/linux/projects/ppc/>

for information about IBM's participation in the two mentioned projects.

In general, if your RS/6000 machine has common hardware reference platform (chrp) architecture, or if your machine is a pSeries machine, it will be supported by the current LinuxPPC distributions. For IBM pSeries current hardware configuration details and supported Linux distributions, see the *Facts and Features* brochures from IBM at this URL:

<http://www.ibm.com/servers/eserver/pseries/hardware/factsfeatures.html>

Note: Older IBM RS/6000 hardware based on the “rs6k” (Micro channel) architecture and the PowerPC Reference Platform® (PReP) architecture is not currently supported by major new distributions. For information on older PowerPC kernels and distributions, you can begin by searching through the [linuxiso.org](http://www.linuxiso.org) Web site at this URL:

<http://www.linuxiso.org>

You can also use an Internet search engine such as Google to search for your particular machine model and Linux. For example, for information about Linux running on an RS/6000 43P model 140 (machine type 7043-140), we keyed in this phrase in the Google search field: 7043-140 && linux and the Google search engine located more than 40 references.

There are versions of BSD Unix for the PowerPC platform as well, and BSD is highly compatible with Linux. You can find more information about BSD at the Web site:

<http://www.bsd.org>

6.3 Consolidation

pSeries Linux-enabled servers provide the capability to run several logical partitions that can be made up of combinations of Linux and AIX on a single server. Logical partitioning on the pSeries provides a cost-effective platform for enterprise workloads while supporting open source standards and solutions. This makes the initial software power base a lot broader in the sense that you can incorporate open source software's rich variety of solutions into the system, while preserving the key capabilities of the AIX at the same time.

For example, consider a typical service provider or Web server environment. It is typically architected as a two or three-tier model. In most installations, there are front-end systems (typically thin servers) to handle caching, proxy, domain name server (DNS), etc. There may then be a second tier of small systems to perform Web application serving using WebSphere or competing products in conjunction with an Enterprise Resource Planning (ERP) or Customer Relationship Management (CRM) products such as SAP.

The third tier of servers runs AIX or Linux on a large SMP system to provide the back office and DBMS functions that require high performance and scalability. In many cases, the first and possibly second tiers are running Linux or Windows NT. This results in more servers, and the need for more staff and expensive software to manage multiple platforms. The Linux LPAR configuration with IBM

Software (shown in Figure 6-5) is intended to minimize these complex tasks, and to offer a more robust, secure solution.

This scenario shows all open source applications being used to consolidate what would normally be separate servers into multiple logical partitions.

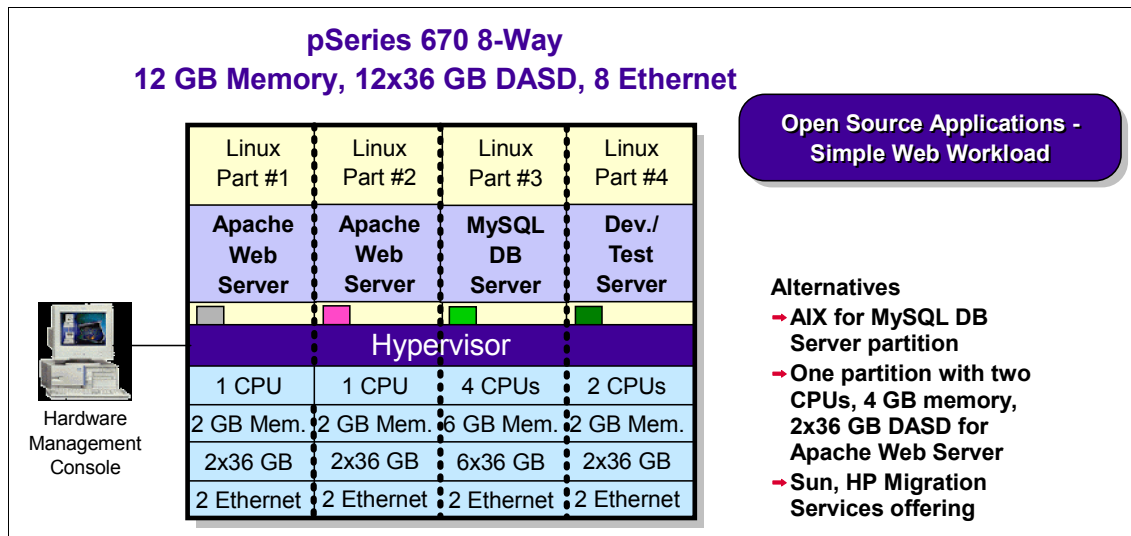


Figure 6-5 Web serving consolidation on logical partitions using Linux for pSeries

This Web serving consolidation can be scaled to a complex Web workload using Software Group products. DB2 can be used in partition 3 to accomplish that. Database management system (DBMS) scalability can be achieved using AIX instead of Linux in partition 3.

6.4 Linux for pSeries clusters

Cluster 1600 is IBM's key AIX and Linux clustering solution for the pSeries. It has increased its scalability and support. Since it also designed to work with CSM, integration of Linux for pSeries-based servers such as pSeries Model 655 is supported. You can learn more about Cluster 1600 on the Web at:

<http://www-1.ibm.com/servers/eserver/clusters/hardware/1600.html>

The POWER4 family of processors provides excellent Linux performance. POWER4 Linux clusters will surely meet present and future market needs. See Chapter 10, "Linux clusters" on page 313, for more information about Linux clusters.

Development of high availability Linux clusters for the distributed pSeries Linux environment is already available and supported.

Beowulf clustering technology and other open source and some commercial products can be used to cluster pSeries systems running Linux to provide compute or high availability clusters. To learn more about Beowulf, see:

<http://www.beowulf.org>

Myricom has the Myrinet switch available for Linux for pSeries. Myricom supplies and supports Myrinet software for Linux 2.4 kernel for IA32, IA64, Alpha, and PowerPC. The Myrinet switch can be used as a high-speed interconnect to cluster systems of pSeries machines running Linux. Gigabit or 10/100 Ethernet connections can also be used. For more information, see:

<http://www.myri.com/scs/>

6.5 Linux for pSeries application solutions

With Linux for pSeries, you can incorporate the boost of open source applications into your environment. Linux for pSeries makes it possible to enrich the capability of any pre-existing pSeries AIX environment with the power and benefits of open source software. IBM and other third-party companies are currently working on providing applications for Linux on the pSeries. For current developments in application solutions for pSeries Linux, see this Web page:

http://www-1.ibm.com/servers/eserver/pseries/linux/sol_all.html

This Web page lists approximately 60 applications (as of this writing), each being a URL link to the particular product's Web page. There is also a convenient filter which lets you list application offerings by category. The current category listings are shown with the number of products in parentheses:

- ▶ Collaborative computing (2)
- ▶ e-business development tools and middle ware (13)
- ▶ Industry applications (18)
- ▶ Infrastructure and system management (17)
- ▶ Scientific and technical computing (9)

The applications range from batch job workload management to compilers, to biomolecular research for cancer.

6.5.1 IBM Software for Linux on pSeries

pSeries servers offer a reliable foundation for Linux-based computing. IBM is working closely with leading Linux distributors to help ensure maximum

performance and functionality. The following Web page, *IBM Middleware Support for Linux*, lists the product support matrix for Linux on all IBM platforms.

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

IBM Java Developer Kit (JDK) Linux

This JDK is a development kit and runtime environment that contains the IBM just-in-time compiler, enhanced with a unique mixed mode interpreter, and a re-engineered Java 2 virtual machine. Like earlier IBM Java versions, the IBM Developer Kit for Linux passes Sun's Java compatibility test and provides stability and performance when you deploy your enterprise e-business applications. For more information, see:

<http://www-106.ibm.com/developerworks/java/jdk/linux140/>

IBM JDK for Linux has been tested with United Linux, SLES7, Red Hat 7, and other distributions for pSeries in a 32-bit and 64 bit environments. For details about the tested Linux distributions for the IBM JDK Linux, refer to:

<http://www-106.ibm.com/developerworks/java/jdk/linux/tested.html>

On December 2, 2003, Oracle issued a news release announcing the availability (Dec. 1, 2003) of the Oracle9i Release 2 (9.2.0.4) Database Developer's Release for Power/Linux. This product release was built on SLES8 (SUSE LINUX Enterprise Server 8) and is capable of running on Linux for IBM iSeries and pSeries servers. Testing with Red Hat Advanced Server 3.0 is in progress. More details about this announcement from Oracle can be found at:

<http://www.oracle.com/start/ibm/powerlinux/intro.html?src=1952614&Act=4>

IBM supports WebSphere Application Server, DB2, and many Tivoli products and development tools on Linux for pSeries. For the latest updates and new software offerings, you can download the *IBM Middleware Support for Linux* document as a PDF file from this Web site:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

This document contains a product matrix of middleware products, organized by platform, for IBM platforms (iSeries, pSeries, xSeries, and zSeries). Each entry in the matrix has the product name, version number, platform, Linux kernel and distribution, additional information sources, and one or more Web links directing you to additional information for that product.

A partial list of some of the IBM products listed in this matrix includes:

- ▶ DB2
 - DB2 Universal Database Enterprise Server Edition

- DB2 Universal Database Workgroup Server Edition
- DB2 Connect editions
- DB2 Universal Developer's Edition
- ▶ IBM Developer and other middleware
 - Cluster Systems Management for Linux
 - IBM Developer Kit for Linux
 - IBM Engineering Scientific Subroutine Library (ESSL)
 - IBM Parallel Engineering Scientific Subroutine Library (PESSL)
 - IBM Visual Age C++
 - IBM XL Fortran
- ▶ Tivoli
 - Tivoli Access Manager
 - Tivoli Configuration Manager
 - Tivoli Directory Server
 - Tivoli Directory Integrator
 - Tivoli Enterprise Console®
 - Tivoli Management Framework
 - Tivoli Monitoring
 - Tivoli Risk Manager
 - Tivoli Storage managers and SAN managers
 - Tivoli Workload Scheduler
- ▶ WebSphere
 - WebSphere Application Server
 - WebSphere Commerce editions
 - WebSphere MQ

6.5.2 Other Linux for pSeries software applications

This section provides examples of pSeries Linux application support in 2003:

- ▶ Absoft: Pro Fortran for POWER Linux
<http://www.absoft.com/newppcproductpage.html>
- ▶ Droplets, Inc.: User Interface V2.2
<http://www.droplets.com/product/>
- ▶ Micro Software Designs LTD: CUMIS Plus
<http://msd-tt.com/products.php3>
- ▶ Platform Computing: LSF 5.1 for distributed and grid computing solutions
<http://www.platform.com/>
- ▶ Point Solutions LTD: PAC Server

<http://www.point-solutions.co.uk/>

- ▶ Samba: Samba File and Print Server

<http://www.samba.org>

- ▶ Myricom provides and supports software for the PPC Linux Kernel to do Linux for pSeries clustering. The Myrinet switch is compatible with Linux for pSeries. See Chapter 10, “Linux clusters” on page 313 for more details. You can also refer to the following Web site:

<http://www.myri.com/>

- ▶ The Globus Toolkit for Grid development

<http://www.globus.org/>

IBM and business partner services solutions are addressed in a separate chapter of this handbook.



iSeries and Linux

IBM has invested to support Linux on iSeries to capture a new generation of applications. Linux enhances iSeries flexibility by enabling another application environment. Linux is a way for iSeries to capitalize on the open source movement and to leverage Linux's virtual world-wide development team. Linux encourages a broad skill base to deliver iSeries based solutions. iSeries can also leverage other IBM hardware, software, and services investments in Linux.

This chapter provides a brief overview of the Linux solution available on the iSeries server. It includes information about the following topics:

- ▶ How Linux runs on the iSeries server
- ▶ Why you should use Linux on the iSeries server
- ▶ Linux and its relation to logical partitions (LPAR) on the iSeries server
- ▶ iSeries Linux integration
- ▶ Linux support on specific iSeries models
- ▶ Linux distributions supported on iSeries
- ▶ iSeries customers and Linux
- ▶ iSeries and Linux as a server consolidation environment
- ▶ Linux application solutions available on iSeries
- ▶ Linux on iSeries service solutions

For additional and current information about the iSeries server and Linux, go to:

<http://www.ibm.com/servers/eserver/iseries/linux/>

7.1 How Linux runs on the iSeries server

The iSeries server provides the ability to run Linux in a logical partitioned environment (LPAR). At present the primary partition must be running OS/400 V5R1 or later, which provides the support required to boot the Linux kernel in a secondary partition. The Linux operating system needs to be installed and run in a secondary partition.

Partitioning on the iSeries server supports multiple images of Linux and OS/400. It is available across the whole iSeries product line. For example, on a one-way iSeries server, up to nine partitions can be supported. On a 32-way iSeries server, up to 31 Linux partitions can be supported. You can move processor and memory resources independently between partitions with a granularity of one hundredth (1/100) of a processor and 1 MB of memory. Shared processor resources can be dynamically moved between Linux and OS/400 partitions. In Figure 7-1 we show a Model 825 that has been split into five LPARs. An OS/400 primary that could be a managing partition, with very few resources and used to control the other partitions. There are also four production partitions two OS/400 and two Linux. In our example these could be an existing OS/400 application with a Linux firewall, a WebSphere application, and an OS/400 development environment all in one package.

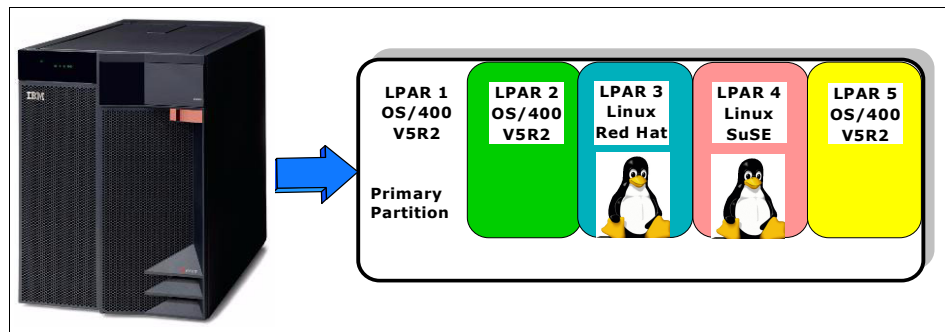


Figure 7-1 iSeries with multiple partitions, two running Linux

Beyond consolidating Linux servers, the iSeries server enables OS/400 applications to be extended with Linux. Linux applications can access DB2 Universal Database (UDB) as well as OS/400 programs and services through one of 16 high-speed virtual Ethernet LANs that connect the partitions. The iSeries server provides integration between Linux applications and OS/400 facilities. Linux applications can access data stored in DB2 UDB in OS/400 and can access files stored in the OS/400 integrated file system (IFS).

The ability to integrate Linux and OS/400 applications on the iSeries server draws together the primary strengths of both operating environments. For

example, the database capabilities of OS/400 and the Web serving qualities of Linux provide a powerful, consolidated computing environment.

7.1.1 Why Linux on iSeries

You may choose to implement Linux on iSeries for the following reasons:

- ▶ **Total cost of ownership (TCO) or total cost of acquisition:** With shared processor and virtual I/O support, there is a minimal hardware requirement for a Linux environment on the iSeries server. There are also space and power consumption savings. The iSeries server can consolidate Linux and Windows NT server farms. This can translate into a minimal dollar investment to run Linux applications on the iSeries server.

iSeries solutions are designed to deliver cost advantages in IT environments. CIOview Corporation provides a tool to assess TCO. With the tool, TCONow!, you can estimate the TCO for selected iSeries solutions compared to competitive alternatives. To use the tool, go to:

<http://www.ibm.com/servers/solutions/serverconsolidation/tco/>

- ▶ **Consolidation:** The iSeries server can consolidate several servers depending on the iSeries model. At the high end of the product line, iSeries supports up to 31 Linux partitions. At the low end, up to nine Linux partitions are supported on a 1-way server. The iSeries server also supports up to 32 OS/400 logical partitions, up to 48 Integrated xSeries servers, and up to 60 connected xSeries servers for Windows Server integration.

Examples of Linux consolidation scenarios in an LPAR environment include:

- *Workload consolidation:* Partition the iSeries server to support Linux workloads that normally require separate servers. Such applications as Enterprise Resource Planning (ERP), Web servers, and business intelligence are excellent candidates.
- *Foot print consolidation:* Partition the iSeries sever to support Linux workloads from different parts of the world.
- *Production and test:* Partition the iSeries server to support Linux production and test workloads.
- *High availability:* Partition and cluster the iSeries server for a Linux high availability solution.
- *Service Bureau, Application Service Provider, Business Recovery Service:* Partition the iSeries server to support a different Linux-based business in each partition.

iSeries provides customers with the ultimate flexibility to allow:

- ▶ **Capacity Upgrade on Demand (CUoD):** iSeries processors can be activated temporarily or permanently, and dynamically allocated to Linux partitions.
- ▶ **Integration:** The iSeries server provides integration between Linux applications and OS/400 facilities. Linux applications can access data stored in DB2 UDB in OS/400 and files stored in the OS/400 IFS.
- ▶ **Performance:** A single large iSeries server with multiple Linux environments can deliver a lower cost per user than multiple systems.
- ▶ **Application flexibility:** The iSeries server supports Java, Domino, UNIX, C, COBOL, RPG, WebSphere, and Windows server applications through xSeries connected servers. Also, with unique virtual disk support, the iSeries server provides an excellent Linux test or development environment.
- ▶ **Resource sharing:** The iSeries server can share resources between OS/400 and Linux partitions. This can reduce the cost of managing and supporting a heterogeneous environment. LPAR supports sharing and dynamically moving processor resources between OS/400 and Linux.

iSeries offers Virtual Ethernet LAN facilities to support partition-to-partition communication. With Virtual Ethernet, you can define up to 16 networks for application communication. These connections can provide 1 Gb performance and do not require LAN adapters, switches, or physical networks.

The iSeries provides Storage Area Network (SAN) facilities for Linux partitions. Linux partitions can access disk, tape, CD-ROM, and DVD resources in OS/400 partitions. With virtual I/O, Linux can leverage the performance availability and manageability of the advanced iSeries storage architecture. iSeries can dynamically add, move, or delete disk space assigned to Linux partitions. The storage spaces for Linux are carved from the OS/400 IFS. This centralization of storage resources enables customers to extend their OS/400 disk subsystem, management skills, and best practices to Linux.

For example, a full OS/400 system backup includes the virtual storage spaces used by the Linux partitions. Daily backups performed by Linux support file-level save/restore and can use the tape devices owned by OS/400.

With shared processor and virtual I/O support, there is a minimal hardware requirement for a Linux environment. This translates into a minimal dollar investment to run Linux applications on the iSeries server.

- ▶ **Reliability:** The iSeries server is a very reliable server and can provide industry leading single server availability.
- ▶ **Fault or failure isolation:** With the iSeries server LPAR technology, Linux workloads or applications can be isolated and result in easier fault isolation, less down time, and faster recovery.

- **Ease of use and systems management:** Linux running in an iSeries server environment can be easier to manage than multiple systems. The iSeries server provides a fast pipe between hardware resources and applications. All of the standard systems management functions such as security, backup and recovery, change and configuration management, network management, and workload management, can be more difficult and time consuming in a multi-site, multisystem environment.
- **Security:** The OS/400 operating environment is a secure operating environment and offers an industry leading secure environment.

iSeries Linux information listed in Table 7-1, are available for iSeries customers. The Web links take you to platform-based information about the iSeries and Linux.

Table 7-1 iSeries Linux information matrix

Information	Location on the Web
Linux for iSeries applications	http://www.ibm.com/servers/eserver/iseries/linux/apps.html
Linux for iSeries Education	http://www.ibm.com/servers/eserver/iseries/linux/educ.html
Linux for iSeries planning	http://www.ibm.com/servers/eserver/iseries/linux/plan.html
Linux for iSeries Developer Resources	http://www.ibm.com/servers/eserver/iseries/linux/resc.html
Linux for iSeries Reference Library	http://www.ibm.com/servers/eserver/iseries/linux/library.html
Linux for the iSeries support	http://www.ibm.com/servers/eserver/iseries/linux/support.html
Linux on iSeries Distributions	http://www.ibm.com/servers/eserver/iseries/linux/dist.html
Linux for iSeries Downloads	http://www.ibm.com/servers/eserver/iseries/linux/downloads.html
iSeries Access for Linux	http://www-1.ibm.com/servers/eserver/iseries/access/linux/

7.1.2 Linux and LPAR

With OS/400 V5R1 and later versions, IBM has enabled Linux to run in multiple secondary partitions. Each secondary partition has its own Linux environment, and can support different Linux distributions. An OS/400 V5R1 or later version is

required to run in the primary or managing partition. The primary partition is used to create partitions and move resources between partitions.

OS/400 V5R2

In OS/400 V5R2, LPAR was enhanced for Linux. With a V5R2 primary partition, shared processor resources can be dynamically moved between Linux and OS/400 V5R2 partitions. In addition, the number of Linux partitions has increased on iSeries servers with SSTAR and POWER4 processors, for example:

- ▶ On a one-processor server, up to nine Linux partitions are supported.
- ▶ On a two-processor server, up to 19 Linux partitions are supported.
- ▶ On a four-processor server, up to 31 Linux partitions are supported.

For V5R2, Linux is supported on the new Model 800, 810, 825, 870, and 890 servers. The 890 supports up to 32 processors and 31 Linux partitions.

iSeries storage virtualization is also enhanced in V5R2. Storage spaces that are allocated from OS/400 can be dynamically added to Linux partitions. Storage virtualization enables Linux to leverage the advanced iSeries storage architecture.

iSeries Navigator is enhanced in V5R2 to support Linux partition management and resource movement. A wizard is available that makes creating a Linux partition even easier.

In January 2003, IBM announced additional enhancements for Linux. With the iSeries Models 870 and 890 Enterprise Edition, an extra processor is activated for Linux at no additional charge. In addition, education and services vouchers are included in the Model 825, 870, and 890 Enterprise Edition servers, which can be redeemed for the Linux on iSeries implementation class and the Linux on iSeries Integration QuickStart service.

OS/400 V5R1

OS/400 V5R1 introduced a new concept for logical partitioning, which allows partitions to share processors. OS/400 V5R1 supports up to four partitions per processor. This opens up the uni-processor iSeries Model 270 and 820 to be partitioned into four partitions.

In addition, a two-way iSeries Model 270 or 820 can support up to eight partitions (2 x 4), and a four-way iSeries Model 820 can support up to 16 partitions (4x4). iSeries servers with eight or more processors support a maximum of 32 partitions. On select iSeries Models 270, 820, and 840 with SSTAR processors, these shared processor partitions can run OS/400 V5R1 or Linux. On iSeries

Models 820, 830, and 840 with ISTAR processors, Linux requires a dedicated processor.

With LPAR on OS/400 V5R1, the minimum processor allocation requirement to create a new partition is one-tenth (1/10) of a processor. This allows the primary partition processing requirements to be reduced to a minimum with the remaining processing units used to support secondary partitions.

An advantage of LPAR is resource movement among partitions. Processor, memory, and I/O resources can be moved among OS/400 and Linux partitions. On select iSeries models, one-hundredth (1/100) of a processor can be moved between partitions. On iSeries models that support only dedicated processors, processor movement among partitions is supported in whole processor increments. With OS/400 V5R1, moving resources to or from a Linux partition requires the Linux partition to be restarted.

A Linux environment on iSeries can be set up simply and inexpensively. The minimum resource for a Linux partition is one-tenth (1/10) of a processor on iSeries servers that support Linux shared processors. In addition, a minimum of 64 MB of memory is needed for Linux. The largest Linux partition is determined by the support provided in the Linux distribution. The 2.4 64-bit kernel supports up to 24 processors and up to 256 GB of memory.

Linux I/O flexibility

The Linux environment on iSeries supports two types of I/O, virtual and direct. With *virtual I/O*, such I/O resources as disk, tape, and CD-ROM are owned by an OS/400 partition. OS/400 shares the resources with Linux. The partition sharing the resources is called the *host*. The host can be the primary partition, but it does not have to be. One partition can be the host for multiple Linux partitions. The shared resources are managed by OS/400. For example, disk space is allocated from OS/400's disk drives for Linux to use. OS/400 provides Redundant Array of Independent Disks (RAID) protection, and some backup and restore facilities for the Linux environment.

The Virtual Ethernet LAN support provided in OS/400 V5R1 and V5R2 extends to Linux. Virtual Ethernet provides 16 1 Gb Ethernet communication paths between partitions (OS/400 to Linux and Linux to Linux) without requiring additional hardware resources.

With *direct I/O*, the I/O resources are owned by Linux. For example, disk drives, LAN, or WAN adapters are under the control of one Linux environment. OS/400 cannot use these resources. A Linux server in a partition can take advantage of both virtual and direct I/O.

7.1.3 iSeries Linux integration

The iSeries server offers many points of integration that support Linux applications leveraging OS/400 applications and data.

A PowerPC Linux Open Database Connectivity (ODBC) driver provides access from a Linux application running on iSeries to the OS/400 DB2 UDB for iSeries database. This ODBC driver is included in select Linux distributions.

IBM has enhanced OS/400 NetServer™ with Samba support for the exchange of files between OS/400 and Linux systems. With this support, Linux can print to OS/400 print queues.

The IBM Developer Kit for Linux, Java 2 Technology Edition, is available for iSeries Linux. This Java Virtual Machine (JVM) in combination with the iSeries Toolbox for Java or its OpenSource version, JTOpen, allows developers to access DB2 UDB for iSeries data through Java Database Connectivity (JDBC) and leverage OS/400 programs and services. The Developer Kit and JTOpen are included in select Linux distributions.

7.1.4 Linux on iSeries models

The iSeries server supports Linux on models that support LPAR and OS/400 V5R1 or later. Table 7-2 shows the iSeries server models by processor feature, whether it supports LPAR, and the number of Linux partitions that are supported.

Table 7-2 The latest iSeries models

Model 800	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2463	V5R2	Yes	Yes	Yes	V5R2	1
	2464	V5R2	Yes	Yes	Yes	V5R2	1
Model 810	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2466	V5R2	Yes	Yes	Yes	V5R2	1
	2467	V5R2	Yes	Yes	Yes	V5R2	1
	2469	V5R2	Yes	Yes	Yes	V5R2	2
Model 825	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2473	V5R2	Yes	Yes	Yes	V5R2	3/6
Model 870	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2486	V5R2	Yes	Yes	Yes	V5R2	8/16
Model 890	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	197	V5R2	Yes	Yes	Yes	V5R2	24
	198	V5R2	Yes	Yes	Yes	V5R2	32
	2487	V5R2	Yes	Yes	Yes	V5R2	16/24
	2488	V5R2	Yes	Yes	Yes	V5R2	24/32
	2497	V5R2	Yes	Yes	Yes	V5R2	16/24
	2498	V5R2	Yes	Yes	Yes	V5R2	24/32

Table 7-3 Model 270

Model 270	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2248	-	No	No	-	-	1
	2250	-	No	No	-	-	1
	2252	-	No	No	-	-	1
	2253	-	No	No	-	-	2
	2301	-	No	No	-	-	1
	2302	-	No	No	-	-	1
	2422 DSD	-	No	No	-	-	1
	2423 DSD	-	No	No	-	-	1
	2424 DSD	-	No	No	-	-	2
	2431	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2432	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2434	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2452Â DSD	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2454 DSD	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
(1) Indicates that dedicated processors are required.							

Table 7-4 Model 820

Model 820	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	150	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	151	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
	152	V5R2	Yes	Yes	Yes	V5R1	4
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2395	V5R2	Yes	No	-	V5R1	1
		V5R1	Yes	No	-	V5R1	
	2396	V5R2	Yes	No	-	V5R1	1
		V5R1	Yes	No	-	V5R1	
	2397	V5R2	Yes	Yes	No	V5R1	2
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2398	V5R2	Yes	Yes	No	V5R1	4
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2425 DSD	V5R2	Yes	No	-	V5R1	1
		V5R1	Yes	No	-	V5R1	
	2426 DSD	V5R2	Yes	Yes	No	V5R1	2
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2427 DSD	V5R2	Yes	Yes	No	V5R1	4
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2435	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2436	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2437	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2438	V5R2	Yes	Yes	Yes	V5R1	4
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2456Å DSD	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2457 DSD	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2458 DSD	V5R2	Yes	Yes	Yes	V5R1	4
		V5R1	Yes	Yes	Yes	V4R5(1)	
(1) Indicates that dedicated processors are required.							

Table 7-5 Model 830

Model 830	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	153	V5R2	Yes	Yes	No	V5R1	8
		V5R1	Yes	Yes	No	V5R1	
	2349	V5R2	Yes	Yes	No	V5R1	4/8
		V5R1	Yes	Yes	No	V5R1	
	2351	V5R2	Yes	Yes	No	V5R1	1/8
	Ã	V5R1	Yes	Yes	No	V5R1	
	2400	V5R2	Yes	Yes	No	V5R1	2
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2402	V5R2	Yes	Yes	No	V5R1	4
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2403	V5R2	Yes	Yes	No	V5R1	8
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	(1) Indicates that dedicated processors are required.						

Table 7-6 Model 840

Model 840	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	158	V5R2	Yes	Yes	Yes	V5R1	12
		V5R1	Yes	Yes	Yes	V5R1	
	159	V5R2	Yes	Yes	Yes	V5R1	24
	Å	V5R1	Yes	Yes	Yes	V5R1	
	2352	V5R2	Yes	Yes	Yes	V5R1	8/12
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2353	V5R2	Yes	Yes	Yes	V5R1	12/18
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2354	V5R2	Yes	Yes	Yes	V5R1	18/24
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2416	V5R2	Yes	Yes	No	V5R1	8/12
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2417	V5R2	Yes	Yes	No	V5R1	12/18
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2418	V5R2	Yes	Yes	No	V5R1	12
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2419	V5R2	Yes	Yes	No	V5R1	18/24
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2420	V5R2	Yes	Yes	No	V5R1	24
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2461	V5R2	Yes	Yes	Yes	V5R1	24
		V5R1	Yes	Yes	Yes	V4R5(1)	
(1) Indicates that dedicated processors are required.							

7.1.5 iSeries Linux distributions

Linux for the iSeries server is available from leading Linux distributors Conectiva, Red Hat, SuSE, and Turbolinux. Turbolinux is a distribution for China and Japan and Conectiva is for Latin America. These Linux distributions are based on Version 2.4 of the Linux kernel for PowerPC. Starting in January 2002, distributions became available for iSeries based on the 64-bit PowerPC kernel. The latest Linux distributions support 64-bit and 32-bit applications. The Linux distributions include the device drivers and other support required to run on an iSeries server:

<http://www.ibm.com/servers/eserver/iseries/linux/dist.html>

Beyond the kernel, these Linux distributions include the operating system, middleware, and applications. For example, they include the Apache, Tux, and Zope Web servers; the sendmail, Postfix, and exim mail facilities and servers; the Samba and NFS file servers; the lprng and cups print servers; and the TCP wrappers and iptables firewalls. The selected distributions also contain key integration technology with OS/400 including ODBC and JDBC support to access DB2 UDB in OS/400.

The current distributions are:

- ▶ Red Hat Enterprise Linux 3, supports 64-bit and 32-bit applications
- ▶ SuSE Linux Enterprise Server 8, supports 64-bit and 32-bit applications
- ▶ Turbolinux Enterprise Server 8, supports 64-bit and 32-bit applications
- ▶ Conectiva Enterprise Server 8, supports 64-bit and 32-bit applications

	SUSE Turbolinux Conectiva	Red Hat
Product version	SLES 8	Enterprise Linux 3 AS
Service Pack		3 Base
Kernel type	64-BIT	64-BIT
Kernel version	2.4.21	2.4.21
Languages	English, German, French, Italian, Spanish, Portuguese (Brazilian), Simplified Chinese, Japanese	Simplified Chinese, Traditional Chinese, Czech, Danish, Dutch, English, French, German, Icelandic, Italian, Japanese, Korean, Norwegian, Portuguese, Brazilian Portuguese, Russian, Spanish, Swedish, Turkish.
GLIBC Version	2.2.5 -139	2.3.2
GCC	3.2.2-35	3.2.3
IBM ODBC Driver for Linux	5.1.0.0.14-3	5.1.0.-0.16
JT Open (Toolbox for Java)	3.1-24	4.1 via download
IBM Developer Kit for Linux, Java 2 Technology Edition	1.3.1-30	JDK: 1.4.1 SR1 via download
ibmsis Linux Storage device driver Supported SCSI controllers - 2748, 2763, 2778, 2757, 2782, 5702	1.19.13a-0	1.2 via download
Journal File System	Vers 1.0.24-7	JFS: 1.1.4 via download
icom serial device drivers	1.1.1	1.1.1

Figure 7-2 Linux distribution support for iSeries

7.2 Consolidation

With iSeries LPAR technology and support for multiple application environments, e-business technologies, and operating systems, you have the flexibility to choose the best applications for your needs. In a single server, you can run your core business and e-business applications such as Lotus Domino on OS/400 or Linux e-mail servers, Web servers, or firewalls. You can centrally manage them while they are running completely independent of one another.

The iSeries server can host and integrate objects and data from another operating system with its own core operating system, OS/400, making it versatile for mixed IT environments. The iSeries supports these operating environments:

- ▶ OS/400 with up to 31 secondary partitions (run OS/400 natively)
- ▶ Linux for PowerPC running natively within an iSeries logical partition. Linux applications can supplement OS/400 applications. You can run up to 31 separate Linux server partitions, all on the same iSeries server.
- ▶ Microsoft Windows servers through the Integrated xSeries Adapter and the Integrated xSeries Server.
- ▶ UNIX through the OS/400 Portable Application Solutions Environment (OS/400 PASE). OS/400 PASE is an environment that is both independent and tightly integrated with OS/400. UNIX applications can be ported to and executed in the OS/400 PASE environment. They can also access objects in OS/400, all with little or no modifications to the original UNIX application.

Figure 7-3 shows how the iSeries works as a consolidation server.

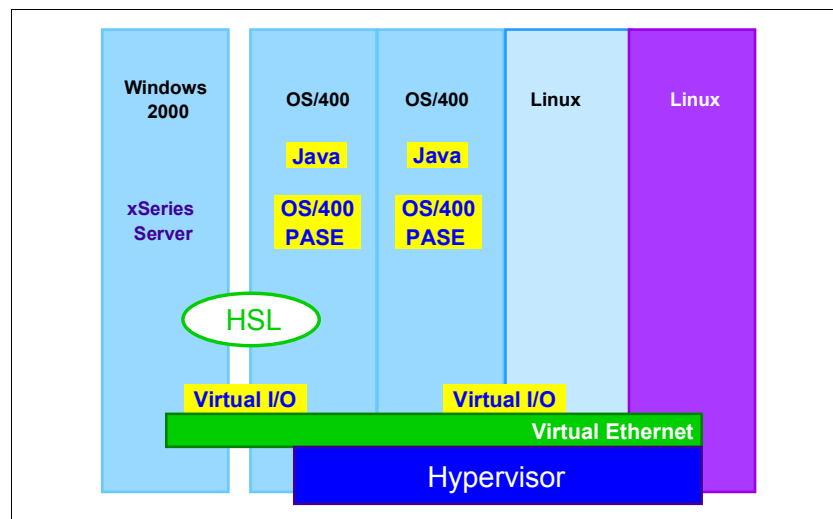


Figure 7-3 iSeries as a consolidation server

7.3 Linux clusters

You can create Linux clusters on the iSeries server. This is possible by using such open source tools as Open Source Cluster Application Resources (OSCAR).

An example solution for Linux clusters on iSeries is H.A. Technical Solutions' EchoStream Data Replication Software, which can automatically replicate data in real-time clusters for Linux. H.A. Technical Solutions - H.A. Clusters product provides a low cost easy to use monitoring system to detect failures and kick off recovery.

Note: One of the major iSeries High Availability Business Partners, Lakeview Technology Inc. of Oakbrook Terrace, IL, recently acquired H.A. Technical Solutions to expand its information availability solutions.

Lakeview Technology Acquires H.A. Technical Solutions, LLC Assets Expanding its Availability Offerings to Linux and UNIX

Visit the Lakeview Technology Web site at for more details:

http://www.lakeviewtech.com/pressroom/pr_2003_11_04.asp

For additional information about Linux clustering technology and solutions, see Chapter 10, "Linux clusters" on page 313.

7.4 Performance

Linux on iSeries workloads are now available in the IBM @server Workload Estimator. Specifically, Samba file serving, Apache Web, WAS, DB2, Network Infrastructure (Firewall, DNS, DHCP) serving workloads can now be defined. Workload Estimator estimates the size of the Linux partition that is needed to support these workloads.

You can access the Workload Estimator on the Web at:

<http://www.ibm.com/servlet/EstimatorServlet>

In addition, a Workload Estimator plug-in is now available for Bynari software. Bynari software is a simple low-cost e-mail alternative to Microsoft Exchange e-mail server and for eOne Group. You enter the mail serving workload characteristics (such as number of mail users and the number of messages sent per hour) into the plug-in. Then, the Workload Estimator tool launches and estimates the size of the Linux partition needed to support the mail workload. You can find this plug-in on the Web at:

<http://bynari.net/index.php?id=1314>

7.5 Linux in a distributed enterprise

Running Linux on the iSeries server combines the strengths of Linux and OS/400 for an integrated solution. Linux delivers excellent open source solutions, while OS/400 is a premier integrated platform for business solutions. Linux enables a new stream of e-business applications for the iSeries platform that complements its strength as an integrated core business solution. Linux applications benefit from the iSeries platform's ability to provide resource flexibility, reliability, and connectivity to other applications on a single server.

When deploying Linux on an iSeries server, you can use the same base configuration each time. Linux is run in an LPAR, so you always need an OS/400 primary partition. When you have a hosted Linux partition, the primary partition can be the hosting partition. Therefore, the base scenario is an iSeries server with a heterogeneous combination of Linux and OS/400 partitions.

It is possible to deliver a reliable all-in-one-box integrated solution to end users in a distributed enterprise. Figure 7-4 shows an example of such an implementation, in which each operating system is used for its intended purpose. OS/400 can be the best platform for business applications. Through the years, it has proven to be quite reliable. Linux on the other hand has proven to be a stable operating system for (inter) networking in general such as providing firewall functions.

In this example, LPAR is leveraged to support numerous operating environments on one iSeries server. Typically, environments such as firewalls and Internet servers are installed on their own servers. Figure 7-4 shows that one iSeries server can support five different servers. It also shows how virtual and direct I/O in a Linux environment can be leveraged. The firewall that interfaces with the Internet has a dedicated LAN adapter. The other resources needed by this partition (disk, tape, CD-ROM, and LAN) are virtual. Virtual LAN is also used to provide a direct connection to the Web-serving partitions.

Two Web-serving partitions running a commerce or business-to-business application can provide a load balancing and high availability environment. The Web-serving partitions are then connected to a second firewall through a different virtual LAN connection. The multiple virtual LAN connections provide another level of security between the outside and inside environments.

The second firewall uses virtual I/O disk, tape, CD-ROM, and LAN resources. It does not have any direct I/O devices except for the native LAN. It is possible to connect the direct LAN adapter to the first OS/400 partition. The advantage of this approach (with the direct LAN adapter connected to the Linux partition) is that all system users must pass the firewall. This means that you can even stop

access to the core business software for some parts of the internal network or log activity to that system.

The OS/400 partition runs the line-of-business applications and contains the data on which the Web applications are based. These five partitions can be set up on a small iSeries server with two processors (for example, a Model 810). The OS/400 and Linux partitions can share these two processors, and resources can be moved between the partitions based on the workload demands. This same solution can be set up on a large iSeries server with additional processors.

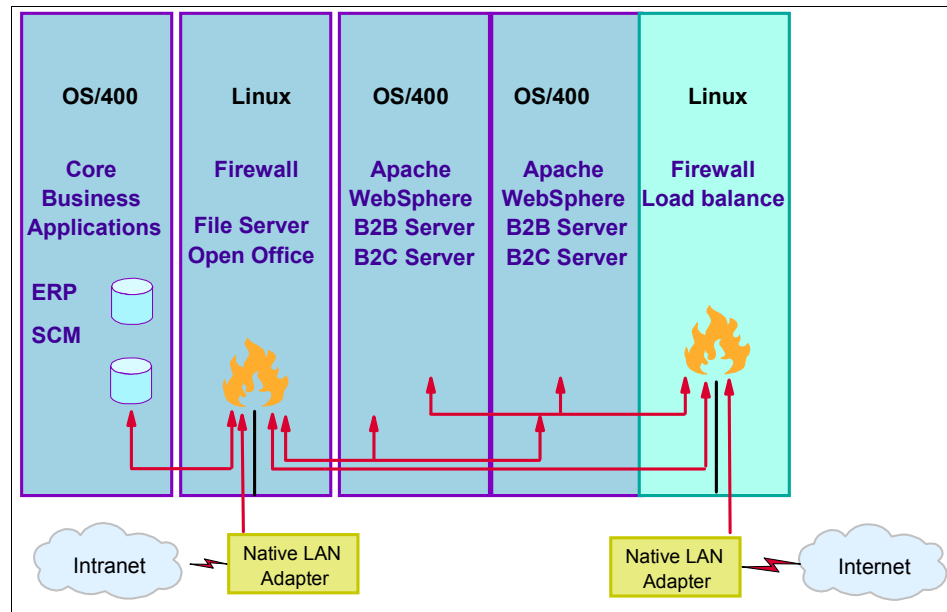


Figure 7-4 iSeries all-in-one integrated solution

7.6 IBM Software Group products for iSeries

IBM Software Group provided support for WebSphere and DB2 for Linux on the iSeries server. For additional information about IBM Software Group products, see 9.1, “Overview of Linux and IBM Software availability” on page 266.

WebSphere for Linux on iSeries

WebSphere Application Server V5.0.2 is now available for Linux on iSeries and pSeries. WebSphere provides an industry leading Web application solution on Linux to support new applications from solution providers and in house development efforts.

WebSphere® Application Server, V5 is the premier Java™ 2 Enterprise Edition (J2EE™) and Web services technology-based application platform, offering one of the first production-ready application servers for the deployment of enterprise Web services solutions for dynamic e-business.

WebSphere Application Server running in a Linux partition on iSeries can leverage the data and applications running in OS/400 to extend applications to new users, customers, and partners.

With WebSphere Application Server for Linux on iSeries, you can leverage the on-demand benefits offered by iSeries including dynamic logical partitioning, shared processor, virtual storage management, virtual Ethernet LAN, and On/Off capacity on demand. You can also leverage the skills and solutions available in the Linux market place. To learn more about the iSeries solution, read the Linux on iSeries White Paper (83 KB):

<http://www.ibm.com/software/webservers/appserv/was/>

WebSphere Application Server products supported:

- ▶ WebSphere Application Server V5.0.2
- ▶ WebSphere Application Server Network Deployment V5.0.2
- ▶ WebSphere Application Server Enterprise V5.0.2

Integrated platform for e-business

If you are evaluating using WebSphere to connect your OS/400 applications and data to new users, suppliers, or partners, the Integrated Platform for e-business can cut the time, skill, and risk associated with setting up an e-business infrastructure. The Integrated Platform is a blue print that sets up three Linux partitions on iSeries: WebSphere Application Server, firewall, and Web server, and provides a connection back to OS/400 and DB2.

DB2 for Linux on iSeries

DB2 UDB V8.1 is now available for Linux on iSeries and pSeries. DB2 provides an industry leading database management solution on Linux to support new applications from solution providers and in house development efforts.

DB2 UDB V8.1 includes DB2 Connect, which provides another option for you to access information in DB2 on OS/400 from a Linux application. IBM also provides migration toolkits for moving data from Oracle, SQL Server, or Sybase.

With DB2 for Linux on iSeries, you can leverage the on-demand benefits offered by iSeries including dynamic logical partitioning, shared processor, virtual storage management, virtual Ethernet LAN, and on/off capacity on demand. You can also leverage the skills and solutions available in the Linux market place. To

learn more about the award winning iSeries solution, read the Linux on iSeries whitepaper (83 KB).

DB2 Universal Database products supported are:

- ▶ DB2 UDB V8.1 Workgroup Server Edition
- ▶ DB2 UDB V8.1 Enterprise Server Edition (and data partitioning feature)
- ▶ DB2 UDB V8.1 Developers Edition
- ▶ Learn more about the features offered by these editions
- ▶ DB2 Connect V8.1
- ▶ DB2 Support of Microsoft .NET Framework
- ▶ DB2 Everyplace®
- ▶ Speed-start your Linux app

Build your next application on Linux — help yourself to trial IBM middleware and tools for Linux, as well as the free training and tech support you need to get started.

Download DB2 documentation, download DB2 runtime, DB2 administration, or DB2 application development clients:

<http://www.ibm.com/developerworks/offers/linux-speed-start/download-p.html>

Migration support for DB2 on Linux

Find resources on how to switch from Oracle, Sybase, Microsoft SQL Server, and others. DB2 leads the industry in reuse and synergy of existing IT assets, removing inhibitors on scalability, information management, and data access across current, legacy and emerging pervasive platforms. For information on how to switch, visit DB2 Migrate Now!:

<http://www.ibm.com/software/data/db2/migration/>

Get answers to frequently asked questions about DB2 for Linux on iSeries (749 KB):

http://www.ibm.com/servers/eserver/iseries/linux/pdfs/db2_linux.pdf

IBM Tivoli Storage Manager, Version 5.2 Client: Linux for iSeries

Tivoli Storage Manager Client for Linux used with Tivoli Storage Manager Server for OS/400 allows customers to back up data from an iSeries Linux partition to most iSeries attachable tape devices.

Tivoli Storage Manager enhances the function of OS/400 backup and restore by providing automated, managed data archive and retrieval, as well as a way to integrate iSeries backups including Linux partitions on iSeries with the backups of other servers in your enterprise. Specifically, Tivoli Storage Manager Client for Linux on iSeries allows a customer to perform file and directory level backups of their iSeries Linux partitions from OS/400, or any other Server running Tivoli Storage Manager Server.

7.7 iSeries ISV application solutions

PowerPC Linux applications run out of the box on the iSeries server, and run pSeries and other PowerPC platforms. The Linux distributions for iSeries contain such notable open source applications as Apache, Samba, IMAP, OpenLDAP, Perl, SendMail MTA, Squid, and PostgreSQL.

The number of Linux based applications for iSeries and pSeries PowerPC platforms has grown substantially during 2003. In this section we will not cover all of these applications. For the most current list visit:

<http://www.ibm.com/servers/eserver/iseries/linux/apps.html>

Linux applications are sourced from many areas. Intel-based Linux applications need to be recompiled for PowerPC. You can download open source applications that are available on the Internet and compile them for PowerPC. Applications can also be developed for PowerPC. Compilers and languages are available with the Linux for iSeries distributions (for example, C, PHP, and Perl). You can also develop applications on PowerPC workstations or servers, and simply move them to run on the iSeries. Cross-compilers are available that support writing applications on Intel systems, and compiling them to run on PowerPC platforms.

Solution providers have announced products for Linux on the iSeries server. These offerings include products from SAE, S2, Lansa, eOne, and Bynari. To assist solution providers with bringing additional solutions to market, iSeries Test Drive for Linux is available and provides remote access to Linux partitions on iSeries. For additional information, see:

<http://www.iseries.ibm.com/developer/factory/testdrive/index.html>

Infrastructure applications

The following table is an extract of the applications available on the iSeries Linux Web site at:

<http://www-1.ibm.com/servers/eserver/iseries/linux/apps.html>

Solution	Company	Description
Transformation Server	DataMirror Corporation	Transformation Server extends ETL functionality by allowing companies to capture, transform and flow corporate and customer data bi-directionally and in real-time.

Figure 7-5 Infrastructure solutions

<u>HA Clusters for Linux, HA EchoStream Data Replication, HA Load Balancing</u>	H.A. Technical Solutions	EchoStream Data Replication - H.A. EchoStream Data Replication Software automatically replicated data in real time Clusters for Linux - H.A. Clusters a low cost easy to use monitoring system to detect failures and kick off recovery
<u>DocHouse (German)</u>	KUMATronik Software GmbH	Catalog manager
<u>DynaLDAP (Korean)</u>	Linux Korea (formerly LDS Technology)	e-business solution set
<u>DynaMail (Korean)</u>	Linux Korea (formerly LDS Technology)	e-business mail solution
<u>AppManager Suite</u>	NetIQ (formerly PentaSafe Security Technologies, Inc.)	AppManager Suite provides auditing, security management and host-based intrusion detection for Linux on iSeries
<u>PostgreSQL Database Management System</u>	PostgreSQL Open Source	PostgreSQL asynchronous replication software
<u>L4 - Load Balancing</u>	Pumpkin Networks	Capacity provisioning and redistribution for peak-load conditions.
<u>Veracity v3.2</u>	Rocksoft Limited	Veracity, Rocksoft's flagship product, is a data integrity security tool that provides multiplatform networked enterprise-wide security.
<u>Sniper</u>	Winstechnet	Security Solution Set
<u>BIMON (Korean)</u>	LinuxSecurity	Firewall
<u>BLAST</u>	Opensource	Data base search algorithms (Basic Local Alignment Search Tool)
<u>Terasolve</u>	Descisys Limited	Descisys is one of the worlds leading software groups specialising in business intelligence. Terasolve is a new high performance business intelligence data warehousing tool. Because of its high performance, it allows the use of all information from operational systems, instead of fragmented or tailored subsets that were required previously for performance reasons. With specific applications for financial management support, and a world class data warehousing practice, Descisys can truly help you extract information to support decision making.
<u>Visma Business ERP system</u>	Visma Software ASA	Visma Software solutions consist of core systems that focus on financial management. All of the financial management systems include complete routines for accounting, orders, inventory, invoicing and purchasing. Connected systems project management, CRM, payroll administration and eCommerce.

Figure 7-6 Infrastructure solutions

ebusiness applications

The following table is an extract of the applications available on the iSeries Linux Web site at:

<http://www-1.ibm.com/servers/eserver/series/linux/apps.html>

Solution	Company	Description
<u>MOBILEmanager (German)</u>	Adisoft AG	With MOBILEmanager, mobile connection to the company's network has been established through the best possible means.
<u>Foedero DocumentZone</u>	Foedero Technologies, Inc.	Foedero DocumentZone is an all-purpose, web-optimized document processing system. DocumentZone is a major breakthrough in building collaborative business networks through secured document access, exchange and sharing.
<u>WorkDesk 2.5</u>	HancomLinux Inc.	Collaborative Office solution
<u>Desknet</u>	NeoJapan	Office Suite
<u>NetOffice</u>	Net Point Limited	Provides web design and Internet solutions for business.
<u>Contento Quest Release 3</u>	Contento Consulting Limited	Enterprise Web Reporting and Application Integration System
<u>Parallel Sort Engine and ETL Interface</u>	CoSORT/Innovation Routines International, Inc.	CoSORT's popular data sorting and ETL acceleration package now tackles high volume select-sort-join-convert-aggregate-report tasks under Linux on iSeries.
<u>Blaze Advisor v5.0</u>	Fair, Isaac and Company, Inc.	Fair Isaac Blaze Advisor is a complete software system for developing, testing, deploying, and maintaining business rules as a part of any automated business system.
<u>FirstSQL/J Enterprise Server</u>	FirstSQL, Inc.	FirstSQL/J Enterprise Server - allows unlimited client access for scalable client/server database applications - Enterprise class, embedded and otherwise.
<u>FirstSQL/J Professional</u>	FirstSQL, Inc.	FirstSQL/J Professional - for embedded, device and desktop database applications requiring single-user access.
<u>FilePro</u>	fP Technologies, Inc.	filePro is a Relational Database Management Tool, a 4GL and a MenuDriven Application Generator. It is an ALL IN ONE Programming Environment- with all of the tools you need for creating high quality, feature-rich applications.
<u>SHAZAM Econometrics Software Version 9</u>	University Of British Columbia	SHAZAM is a comprehensive computer program for econometricians, statisticians, biometricians, sociometricians, psychometricians, politicometricians and others who use statistical techniques. The SHAZAM command language has great flexibil

Figure 7-7 e-business solutions

<u>Introscope SQL Agent 4.0</u>	Wily Technology Inc	Introscope is an Enterprise Java Application Management solution that helps companies ensure the performance and availability of their mission-critical applications and supporting systems.
<u>Acu4GL for DB2</u>	Acucorp, Inc.	Access to Informix, Oracle, Sybase, and ODBC - compliant databases from applications written in ACUCOBOL-GT using SQL queries.
<u>ACU Connect</u>	Acucorp, Inc.	ACUConnect is a distributed computing technology that lets developers split application processing between multiple machines for optimal performance.
<u>AcuCOBOL-GT Dev System</u>	Acucorp, Inc.	The ACUCOBOL-GT development system enables COBOL developers to implement and deliver full featured GUI-enhanced, Web-enabled COBOL, DB2-integrated applications.
<u>AcuODBC Server</u>	Acucorp, Inc.	AcuODBC communications interface is designed to provide easy access to COBOL data files from popular Windows applications
<u>AcuSQL</u>	Acucorp, Inc.	AcuSQL™ product supports the inclusion of embedded SQL (ESQL) statements in ACUCOBOL-GT source code
<u>AcuServer</u>	Acucorp, Inc.	The AcuServer™ remote file server technology provides file access services to ACUCOBOL-GT applications
<u>CISAM Interface</u>	Acucorp, Inc.	The Acu4G technology provides seamless access to DB2, Informix, Oracle, Microsoft SQL Server, Sybase, and ODBC-compliant databases
<u>ACUCOBOL-GT Runtime</u>	Acucorp, Inc.	ACU Cobol environment on iSeries Linux without recompilation.
<u>Web Hosting Automation Software</u>	Alabanza Corporation	It is the Web hosting solution for Internet service providers, telecommunications carriers, systems integrators, Web designers and developers via its four powerful applications: the Domain System Manager, Control Panel, and Systems Administration.
<u>Info @ Work</u>	Brookline Technologies, Inc.	Serve as a single point of interaction with information, applications, and people; secure, personalized, instant access to applications and knowledge resources; enable self service for employees, partners, or customers.

Figure 7-8 e-business solutions

<u>Insight Server V4</u>	Bynari Inc.	Email and calendar server
<u>WebSales v1.2 Petit EOS Edition for iSeries (Japanese)</u>	ClassCat Co., Ltd.	Web Based Business to business (B2B) solution
<u>Mini WebSales v1.2 Infor EDI Edition for iSeries (Japanese)</u>	ClassCat Co., Ltd.	Web based B2B solution set for Sales.
<u>Dynamic Personal WebShop 1.2 Intro for iSeries (Japanese)</u>	ClassCat Networks, Co., Ltd.	Web creation tool set for B2B environments.
<u>Intertrac</u>	ComputerWorks	Software tool equipped for complete enterprise-wide information management.
<u>Nimbus v2.5</u>	Converse Software LLC	NimBUS is an IT management platform that provides scalable, resilient and reliable monitoring capabilities for organizations that wish to proactively manage critical IT resources.
<u>Infinity Portal</u>	CrossLogic Corporation	Infinity is a feature-rich content delivery engine that can capture information from virtually any source on the web or in an enterprise, and create a customized portal for every user in the organization.
<u>Infinity Commerce Server v2</u>	CrossLogic Corporation	CrossLogic offers a fantastic quickstart solution for getting started with Java and IBM WebSphere.
<u>Strata Frameworks</u>	CrossLogic Corporation	The Strata Frameworks is a streamlined way of creating and implementing object-oriented, well designed Web applications.
<u>Eitasoft Single Sign-On Server (SOS)</u>	CrossLogic Corporation	The Eitasoft Single Sign-On (SSO) Server is a powerful authentication tool that allows users to login once and connect to all systems in an organization. Infinity SSO is perfect for Intranet, Extranet and Portal applications.
<u>e-vendo</u>	e-vendo AG	eBusiness ERP and Marketplace solution set.
<u>Easysoft ODBC-ODBC Bridge</u>	Easysoft Limited	The Easysoft ODBC-ODBC Bridge is a universal ODBC driver that allows Linux applications to access any ODBC compatible data source on a remote system.
<u>Forceten</u>	EEDO Knowledgeware Corporation	ForceTen enables you to provide employees, partners and customers with push and pull access to knowledge via e-learning, on-the-job support, a searchable knowledgebase and knowledge sharing, all driven from a single database.
<u>eOne Commerce Suite</u>	eOneGroup	eOneCommerce is an affordable, highly scalable, (100% Java, XML) e-commerce application for companies of all sizes.
<u>ECommerce (German)</u>	Fritz & Macziol GmbH	Shipment Automation and Shipment Security - Firewall.

Figure 7-9 e-business solutions

<u>CM SAFE(Configuration Manager)</u>	GMx Solutions LLC	Provides robust iSeries Configuration-File Management along with capabilities for Disaster Recovery.
<u>RIITA Server</u>	GO Software	A payment processor middleware for iSeries and othe Linux capable platforms. It is designed for businesses that want to accept credit cards, checks, and other payment forms for their retail point of sal (POS), direct marketing and eCommer
<u>HDE Center (Japanese)</u>	HDE	Application development tool
<u>HiT JAllora XML-RDB</u>	HiT Software, Inc.	Allora, an XML platform for data exchange, is a set of design-time tools and run-time engine that enables applications to perform real-time, bidirectional transformation between XML data and relational data.
<u>HiT JDBC/DB2</u>	HiT Software, Inc.	SQL middleware, including high performance ODBC, OLE DB, .NET and JDBC, enables applications to read and write data to IBM DB2 databases.
<u>CM SAFE (Configuration Manager)</u>	H.A. Technical Solutions	H.A. Clusters™ High Availability Clustering Software from H.A. Technical Solutions is high availability software designed to manage and control the automatic failover of applications running on servers that are networked in a client server.
<u>WebSphere Commerce Business Edition for Linux v5.5</u>	IBM Software Group	IBM WebSphere Commerce Business Edition Provides a powerful, flexible infrastructure based on a unified platform for running large, high-volume B2B and advanced B2C e-commerce Web sites for global e-businesses.
<u>E-Fusion v1.5</u>	iConexio Technologies Inc.	iConexio's E-Fusion integration technology is based on breakthrough software integration adapter concepts, design patterns, and innovations. The E-Fusion product suite consists of an Adapter Development Environment(ADE) and Adapter Runtime.
<u>Idea.linux for KM</u>	IDEA TECHNOLOGIES SRL	Web enablement tools.
<u>Inovis</u>	IPNet Solutions	Business transaction
<u>Edge Server</u>	LANSA	Edge Server on Linux.
<u>Syndeo Project Server</u>	Macadamian Technologies Inc.	Syndeo Project Server is the enterprise portal for streamlining the software development process. Project Server automates project management, enables collaboration, and integrates development tools.
<u>eG-Model</u>	OEC Co.,Ltd	eBusiness solution development tool
<u>Tsunami WebStream application</u>	PDX BizSystem Sdn Bhd	WebStream is a web ERP system. It extends outside the four walls of the enterprise to customers and suppliers. It is a realtime online application system on the Internet. The package focuses on business-to-business functions.

Figure 7-10 e-business solutions

<u>Enterprise Solution Server</u> (German)	Pentaprise GmbH	e-business solution set.
<u>Corporate Website</u> (German)	PixelPark AG	e-business web site design and multimedia solution set.
<u>GraphGenerator</u> (Japanese)	Plat IS Inc	A toolset to help legacy data on iSeries to go to Web.
<u>NART system</u> (Japanese)	Plat IS Inc	B2b (buy-side) server that automatically sends procurement order from legacy LOB application to the supplier's Web site.
<u>WebExtender</u> (Japanese)	Plat IS Inc	A toolset to help legacy data on iSeries to go to Web. Tools include WebPrint (transforms SCS data into PDF on Linux), HTML generator (convert DB2 UDB data into HTML) and GraphGenerator.
<u>Groupware</u> (Japanese)	Plat IS Inc	Japan's top selling GroupWare now runs on Linux collaborating with DB2 UDB.
<u>P@Bamboo</u> (Italian)	P@rtners S.p.A.	eBusiness web enablement solution set.
<u>ResQ/ME</u>	ResQNet.com	Web enablement tools.
<u>ResQPortal</u>	ResQNet.com	Web enablement tools.
<u>ResQNet</u>	ResQNet.com	Web enablement tools.
<u>Retail Pipeline</u>	Retail Pipeline Integration	Retail Resource Planning system
<u>E.ssential</u>	Rippe & Kingston Systems	The E.ssential Enterprise Portal is a self-service solution that integrates vital applications, documents, databases, news feeds, reports, and more for delivery via a web browser to any community of users.
<u>UC4:global</u>	SBB Software GmbH	UC4:global provides the most comprehensive job scheduling functionality for SAP systems and the entire system environment
<u>Aviva Host Integration</u> <u>SDK</u>	Scanpak, Inc.	Scanpak Inc., a leading supplier of automatic data collection devices and services has developed and deployed a cost-effective and reliable application for the tracking and tracing of fixed assets
<u>SCOPELAND 2000 XML-</u> <u>based application server.</u>	Scopeland Software GmbH	Scopeland 2000 (App Dev Environment) enables you to develop you business application much faster and, even better, much more flexible against changing requirements.
<u>'Ecoro Eh 0.97.4</u> (German)	SHD Datentechnik GmbH & Co KG	eCommerce Solution set
<u>StreamServe</u> <u>Communications Server</u> <u>3.01/ Business</u> <u>Communication Platform</u>	StreamServe Ltd.	StreamServe enables organizations to produce and exchange critical documents independent of source data, format or channel: using its pioneering communication hub.
<u>iBase v4.0</u>	Temenos Headquarters SA	Banking software solution set.

Figure 7-11 e-business solutions

<u>SafeMail</u>	Veridis S.A	SafeMail ensures security of messages and attachments in transit over the Internet. It enables a full-circle point-to-point secure transfer of e-mails and files, in combination with value added services such as notarization.
<u>Avatar v4.0</u>	Vexus	Avatar is a simple yet feature rich batch job scheduling utility.
<u>Jxel 4.614b</u>	WBS Inc.	Web server acceleration and monitoring.
<u>Stato2</u>	Abako Media Oy	Stato is a publishing system for professional use.
<u>bizM Sentio</u>	MailBook Co. Ltd.	e-business solution set.
<u>Insight Web Client</u>	Bynari Inc.	The Insight WebClient is a browser-based application that reads information from Insight Server, and displays this information to the user in the form of an html page.
<u>FrontPage 2002 Server Extensions</u>	Ready-To-Run Software, Inc.	Utilizing Ready-to-Run FrontPage Server extensions enables customers who have developed a Web presence with Microsoft's Front Page technology, can now run that code under Linux on iSeries. Utilizing Ready-to-Run Front-page Server extensions on the same server where the data resides can greatly improve performance, and simplify the management of Web applications by reducing the number of servers that need to be administered.

Figure 7-12 e-business solutions

Web-enablement applications

The following table is an extract of the applications available on the iSeries Linux Web site at:

<http://www-1.ibm.com/servers/eserver/series/linux/apps.html>

Solution	Company	Description
<u>360Commerce Back Office</u>	360COMMERCE	The browser-based 360Store Back Office handles essential inventory management and store operations reporting and analysis.
<u>ABAS-EKS</u>	ABAS Software AG	ABAS offers an ERP which covers Production Planning and Control/ERP Ware-house Management/MRP Materials Management, Stocking Management, Scheduling Production, Sales/Sales Order Processing, Purchasing/Purchase Order Processing, Accounting.

Figure 7-13 Web-enablement solutions

<u>Advantage Series Accounting Package</u>	ACCPAC International, Inc.	The ACCPAC Advantage Series solution provides end-to-end business management applications for small- and medium-size businesses including accounting, CRM, e-Commerce, warehouse management, human resource management, and point-of-sale.
<u>Staff Efficiency</u>	ATOSS Software AG	ATOSS offers a software-package for resource allocation, worktime management, process management, project & performance management and knowledge management.
<u>NoStockOut Retail Suite</u>	Brookline Technologies, Inc.	Fully featured e-commerce solution with a large site currently selling over 80,000 products online, integrate inventories from multiple distributors/suppliers, provide links to manufacturer product & warranty information.
<u>ETM Web Pipe</u>	Caminus LLC	<u>Web-based browser application for service providers.</u> <u>Acquired by Sungard</u>
<u>NOW</u>	DATATEX	NOW is the infrastructure part of the Textile ERP Solution. NOW Web based solutions are used to inventory and warehouse management, purchasing and sales.
<u>Docubase Enterprise Suite for Document Imaging</u>	Docubase Systems - sold via DalTech International Inc	DOCUBASE ENTERPRISE is a EDMS (Electronic Document Management System) suite of products. The Docubase solution includes the capture, indexation and filing of paper documents such as office paperwork, graphics, etc.
<u>Cyber Education (Korean)</u>	DUNet	Online education
<u>Accounts Payable</u>	Ellenbogen Software Associates	Accounts Payable generates all required reports, documents and audit trails.
<u>TempServ</u>	Greenfield Software	TempServ is a comprehensive staffing package seamlessly integrating front and back office. Powerful employee skill and availability search, job order tracking (both filled and unfilled), and complete assignment history by employee
<u>Laguna Document Management (Italian)</u>	Gruppo Pro Spa	Document Management.
<u>eLogistIQ</u>	GUS Group AG & Co. KG	eLogist1Q is an online ordering system. Address online ordering functions like payment, security, and customer profiling.
<u>OS Laboratory Information and Management System (LIMS)</u>	GUS Group AG & Co. KG	OS Laboratory information system handles workload control and integrates into email.

Figure 7-14 iSeries Web-enablement solutions

<u>ERP Financials Solution</u>	Hansa Business Solutions (UK) Limited	Hansa Financials for e-Business supports a full range of eCommerce solutions.
<u>Reflex</u>	HARDIS	Reflex is a global logistics management package designed to resolve all warehousing problems.
<u>Javias</u>	Industrial Application Software GmbH (IAS)	JAVIAS is a web based e-ERP system which covers all areas of a company with complete logistic, PPC and finance functionality
<u>Pharma Suite - Team Player with Cruscotti (Norwegian), Progettometro, ISO applications</u>	Interweb	Applications, fully WEB enabled, with features of Content & WorkFlow Management, Portal, Profiling and Publishing of Shared documents. It allows full integration with all the Java enabled technologies and can be expanded.
<u>Pharma Solution suite - SFA@Libero Farma (Norwegian)</u>	Interweb	A complete soluton for managing the needs of the Pharmaceutical Industry. Some of the components are: TeamPlayer for Pharma:customization of TeamPlayer (teamWare, Content & WorkFlow Management and Portal platform by e-Services) for the Ph
<u>Genscreen</u>	Linux Business Accounting Systems	A C/C++ Code Generator which automatically builds user interface and I/O screens based on the ncurses library of Linux. Converts a flat text file depicting a screen design into compilable code. This will be of great interest to programmers
<u>Linux Ledgers</u>	Linux Business Accounting Systems	Linux General Ledger is intended for use in small to medium sized for general bookkeeping
<u>RadiusWireless Authentication Solutions (Korean)</u>	Linux Korea (formerly LDS Technology)	RadiusWireless Authentication Solutions
<u>eMerchant - CRM solution</u>	Magic Software Enterprises, Inc.	Magic eMerchang v2 is a highly customizable e-commerce framework designed to deploy a e-business CRM solution
<u>Mapics Server Front End</u>	MAPICS, Inc.	MAPICS' solutions include two enterprise resource planning (ERP) foundations, plus supply chain management, collaborative commerce and maintenance and calibration management fuctionality.
<u>Food/Warehouse Industry</u>	Merit-Trax Technologies Inc.	Food Industry Supply Chain
<u>CRM</u>	MetaWise, inc	CRM Solution (Korean)
<u>Gas v6.56</u>	NetFuel Inc.	management/monitoring
<u>CRM Call Center - Permission Marketer (Korean)</u>	Nexus Community Inc.	CRM Call Center
<u>Loga 2001, Loga HRMA</u>	Personal & Informatik AG	LOBA 2001 is multilingual, universally applicable payroll processing software. LOBA 2001 incorporates remuneration in salary and wages along with travel expense accounting.

Figure 7-15 Web-enablement solutions

<u>e-Voice Tour</u>	RealPerson, Inc.	e-Voice Tour on iSeries enables media-rich Internet publishing on web sites, sales force presentations and web-based training.
<u>Reddot CPS</u>	RedDot Solutions AG	ECM Suite
<u>Line 500 Software Product</u>	Sage	Business management application, Line 500, allows mid-sized businesses to take a single view of their finance, distribution, manufacturing and service operations.
<u>SelectorPro</u>	Systems Application Engineering, Inc. (SAE)	SAE offers mobile solutions for various industries, including distributors, where employees wear wrist computers to help them pick merchandise for shipment.
<u>Sales Force Automation</u>	Trabas	Complete end-to-end business solutions from technology and e-business applications
<u>UNISOL Job Accounting</u>	UniSolutions Associates	UNISOL JobAcct® provides system and database accounting and chargeback at the user, group, project and cost-center level, generating usage reports for one or more computers on the network.
<u>Taxware Sales and Use Tax Solution</u>	Velosant	Velosant's Taxware Sales and Use Tax Solution is a robust and flexible tax calculation and reporting system, combining market leading tax research and technology. Operating as an integral part of existing financial or accounting systems.
<u>Orion</u>	Vision Solutions, Inc	ORION is an availability solution, designed to manage both data and application availability across an enterprise's heterogeneous environment.
<u>Enterprise Resource Solution</u>	ZOPE Corporation	Enterprise resource Program with a workflow for managing travel, purchase and vacation requests.
<u>P@CRM (Italian)</u>	P@rtners S.p.A.	Customer Relationship Management solution set.

Figure 7-16 Web-enablement solutions

Small office example using OpenOffice

The iSeries server offers a consolidated server for the small office. OS/400 runs the line-of-business applications while Linux runs OpenOffice for personal productivity applications. OpenOffice is an open source product that provides a word processor, spreadsheet, and presentation builder. Based on the popular StarOffice product, OpenOffice can read and write the respective Microsoft Office files. Sun provided StarOffice to open source resulting in OpenOffice.

OpenOffice can run on a server such as the iSeries, and sends the graphical user interface (GUI) to X-Windows clients. This is a great solution for the small office environment. OpenOffice is installed. After it is on the server, users can customize the OpenOffice environment such as the default printer.

OpenOffice is available for all types of client systems. You can download it from the Web at:

<http://www.openoffice.org/>

Java

iSeries supports current levels of Java with the JVM and the IBM Toolkit for Java. For additional information and downloads, see:

<http://www.ibm.com/servers/eserver/iseries/toolbox/>

7.8 iSeries solutions

The iSeries is the industry’s leading business server for mid-market companies. It features the intelligent integration of key e-business functions, innovative technologies, application flexibility, superior performance, and the new tools required for managing e-business. iSeries Linux solutions are available for you regardless of the size or complexity of your business.

This section gives examples of Linux solutions that were architected on the iSeries server.

International customs and trade services group

An international customs and trade services group needed to increase its information system capacity and not increase its staffing.

The project was implemented with an iSeries Model 820, configured with nine Red Hat Linux partitions.





















iSeries 820				
0	GHY820 OS/400			0.50 768 M
1	DOCIMG			0.20 256 M
2	INTRANET			0.10 256 M
3	LANAPPS Samba/DHCP/sendmail/SSH			0.50 256 M
4	FIREWALL Firewall/Squid			0.20 256 M
5	L2TPVPN			0.10 256 M
6	IPSECVPN			0.10 256 M
7	INETAPPS Sendmail Gateway/DNS/SSH			0.30 256 M
8	WEB SRV			0.20 256 M
9	TESTLNX			0.1 256 M

Figure 7-17 Proposed system

Food service distribution

One of the largest food service distributors in the South region of the USA sort to resolve the following issues with their current IT environment:

- ▶ Server upgrades, disaster planning, viruses
- ▶ Backup, recovery, storage, redundancy
- ▶ Decentralized support, cost, user administration
- ▶ Consolidate AS/400 170, NCR, and Intel servers

Their solution was:

- ▶ OS/400 RPG and Domino
- ▶ Linux Unix-based Cobol application
- ▶ Linux Web server
- ▶ Three Integrated xSeries Servers Windows applications: Payroll, SQL Server, printer

Energy company

A diversified energy company including power generation, fuel production, electric utility, and broadband services sort to improve the physical security of their facilities without on-site people and provide an audit trail of movements.

Their current IT solution was two i825s in HA configuration with five OS/400 partitions running customer care, billing, financials, and a Domino server. To run their new security application, they decided to create one Linux partition.

The application drives a software motion detector, keeps images that change with time. An Apache server runs a browser application and users can visit a Web site to see live camera feeds from seven security cameras. To make this happen, the company chose an IBM Global Services offering for LPAR implementation.

Outdoor power products

In Sweden one of the world's leading manufacturers of outdoor power products such as chain saws, lawn mowers, and rototillers decided to look at Linux to solve some business challenges. With more and more sensitive data being exchanged over its online ordering application, this company decided it needed additional security.

The outdoor power products parent company raised the security requirements for all of its systems that were exposed to the Internet. The challenge was to find a cost-effective security solution that could be implemented quickly.

The IT group reconfigured a iSeries 270 system running OS/400 DB2 and WebSphere. They created a third partition running Apache on Linux to separate the Web server from the WebSphere application server.

Chemical distributor

A leading global player in the distribution of chemicals and raw materials. The company sort to resolve problems with 49 server Windows-based network. They wanted a consistent, open, and easy-to-manage infrastructure. They also wanted to give employees the same access to information and services regardless of their location, and provide a more reliable service to the business. This was all to occur against a backdrop of lower costs, by reduction of the number of physical systems.

The solution was Linux-based infrastructure applications on iSeries 825 and xSeries, with a Domino running on OS/400.

Read more about the story at:

<http://www.ibm.com/software/success/cssdb.nsf/CS/DNSD-5QJENP?OpenDocument&Site=default>

Financial services company

A financial services company wanted to optimize their IT Infrastructure. Their installed systems were an AS/400 730, with HP Document Management and infrastructure servers.

The solution involved the company purchasing an iSeries Model 825. They created one OS/400 partition and used it as a portal in the bank, where users can inquire and retrieve indexed images. They also created a Linux partition for the storage of document images (such a scanned checks) in 1 TB of disk. This partition also included a mail server, and open office.

To further their server consolidation, they consolidated their external Intel servers into three integrated xSeries servers. These new servers will run DHCP, the file/print server, and hot spare.

To achieve a satisfactory implementation, the bank chose IBM Global Services to assist with the implementation.

Provincial Education, Korea

This teachers' association decided to improve the integration and management of their research resources such as the documents, photos, and movies managed by the province. They also wanted to reduce costs and improve manageability of the NT server farm.

The Provincial Education created a solution that used a iSeries Model i810 with one OS/400 partition to act as a database and store content. They also chose to implement two Linux partitions, with one running WebSphere Applications to access content. Plus the second to handle their development environment. Their

consolidation with further enhanced with one integrated xSeries Server used for PDF creation.

Power Plant Company, India

The Power Plant division of this Indian company manufactures off-highway automatic and power-shift transmissions.

The business decided they needed to extend its ERP applications to suppliers and large international customers to streamline business processes and reduce costs. Their chosen solution was an iSeries Model 270 to run a partition with Red Hat Linux for e-mail, and an OS/400 partition for ERP and WebSphere for e-business.

They also decided to install an Integrated xSeries Server to run their PC Antivirus application.

Manufacturer of high quality ceramics/china products

A major manufacturer and distributor of high-quality ceramic and china tableware, with a 200-year old craft-led heritage. Their challenge was to develop a robust and scalable commerce for an e-business supply chain management (SCM) system. Formidable overseas competition meant that the manufacturer had a crucial business need to adapt to the world of e-business.

The architected solution for the business was developed using Wizz400, a Clover product, running in a LPAR environment, and a Linux-based firewall system also running in an LPAR. The iSeries server chosen was a powerful iSeries Model 820.

Wizz400 is middleware that is an integral part of the manufacturer's e-business system. Their Web site turned into an interactive e-business tool that can extract live data, display accurate stock and account information, and collect order data.

After transitioning to their iSeries Linux architected solution, the company experienced the following benefits:

- ▶ Decreased costs through consolidation of three servers to one
- ▶ Adaptability that includes multilingual capabilities
- ▶ Adaptable and flexible architecture that allows further development to improve customer service
- ▶ Increased accuracy of orders, stock
- ▶ Improved customer service
- ▶ Improved and simplified administration
- ▶ Increased performance and scalability

To learn more about this story, see:

7.9 iSeries service solutions

Support for Linux on iSeries has two key components. Support for the integration of selected Linux distributions on iSeries is provided by the iSeries support organization. Integration support includes help with installing Linux successfully, creating partitions for Linux, and sharing iSeries resources with the Linux partition. This integration support is provided with an iSeries OS/400 SupportLine contract.

Support for the Linux distribution itself is a separate option. This support is available from SuSE, Turbolinux, Red Hat, IBM Global Services, and other Linux service providers.

7.9.1 iSeries Linux resources

The following are resources available for planning, testing, and implementation of Linux on iSeries.

iSeries Editions

In the Americas the iSeries Enterprise Edition includes a Linux Distribution and one activated processor for Linux use. The Enterprise Edition also includes an education and a services voucher. This is a great method for skills and implementation enablement.

The following link, which discusses the iSeries Editions, and the services and education offerings, is available at:

<http://www-1.ibm.com/servers/eserver/series/hardware/editions/pdf/faq.pdf>

iSeries Technology Center

The iSeries Technology Center (iTC) offers a range of Linux related services, skills transfer, and education in advanced topics on iSeries. Visit the iTC at for more information on what is available:

Education:

<http://www-1.ibm.com/servers/eserver/series/service/itc/educ.htm>

Consulting services:

<http://www-1.ibm.com/servers/eserver/series/service/itc/services.htm>

LOADRUN CD for Linux on iSeries

This is a fast and simple method of producing a Linux image on iSeries. The LOADRUN CD contains a Linux Server Image. This image is configured and tested. It also contains OS/400 CL scripts to assist with the installation.

There are three simple steps:

1. Create the partition.
2. Install.
3. Install and configure the Linux server with a single **LOADRUN** command.
4. The Linux applications are ready to go.



Figure 7-18 Three step Linux install on iSeries

IBM Business Partners can shorten the time required to install and configure Linux to minutes. They can also reduce the skill required to install and configure Linux. This also reduces the risk associated with a Linux installation.

iSeries and Linux redbooks:

Implementing Linux on iSeries, SG24-6232

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246232.pdf>

WebSphere for Linux on iSeries: Implementation Guide, SG24-6958

<http://www.redbooks.ibm.com/redpieces/pdfs/sg246958.pdf>

Bynari Insight Server V4 for Linux, SG24-6069

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246069.pdf>



TotalStorage and Linux

Taking advantage of the open source foundation, Linux-enabled storage solutions free businesses from the boundaries of closed, proprietary systems for impressive savings in hardware, software, and management. In this chapter we discuss IBM TotalStorage solutions on Linux for today's business-critical storage applications, especially in consolidated storage environments.

8.1 Linux and IBM TotalStorage products

Linux is a revolutionary open source platform that is stable, secure, scalable, and powerful, offering today's businesses the flexibility to innovate for success. Storage is a natural component of a well-designed, Linux-based e-business infrastructure. IBM TotalStorage can fortify your Linux investment in addition to helping you cut costs, consolidate infrastructure, and position your organization for the new on demand world, with the strength and reliability you need to manage your data.

Storage Solutions Group is proud to provide connectivity to its entire product lines, hardware, and software. Listed in this chapter are IBM's current Linux offerings.

IBM TotalStorage intends to extend support for the new Red Hat Enterprise Linux version 3 to all its key product offerings, hardware, and software storage products, including the availability of advanced storage functions.

Note: Find more information about Red Hat Enterprise Linux at:
<http://www.redhat.com/software/rhel>

8.1.1 Why storage solutions

Linux is one of the fastest growing, most popular server platforms available today. Analysts' projections for Linux growth range from 28 to 35% compound annual growth rate (CAGR) from 2001 to 2006 compared to other platforms that are expected to grow about 15% during the same period. With its exceptional flexibility, reliability, and scalability, the power of Linux as a business platform is undeniable.

Linux is driving the next generation of e-business networks by giving companies open access to the continuous innovations being provided by the open source community. What is more, Linux offers freedom from the constraints of proprietary systems. This much is clear: Linux is open and ready for business.

Storage consolidation

Consolidation can increase the cost-effectiveness, efficiency, and performance of data storage. With IT administrators spending an ever-increasing proportion of their time as well as their budget addressing storage issues, the advantages of consolidation cannot be ignored.




Figure 8-1 Simple storage consolidation diagram

Storage consolidation unites hardware, software, services, and systems with sound storage management principles. The key objective of storage consolidation is to achieve higher efficiencies and better returns through end-to-end business and infrastructure optimization, making storage an important element to consider when evaluating workload consolidation.

Storage consolidation goes a long way toward improving the manageability and scalability of storage systems. By simplifying storage infrastructures, consolidation makes it easier for administrators to manage growth, respond quickly to changing business needs, and reduce the amount of time spent on deployment and configuration. In addition, thanks to technologies such as redundant array of independent disks (RAID) and clustering, consolidation practices help deliver highly available and fault-tolerant storage for continuous and reliable data access, sharing, and backup.

Linux-enabled storage solutions support impressive business efficiencies in the form of low TCO and high ROI; increased openness and flexibility for highly interoperable storage solutions; and the reliability, scalability and security required by vital data systems. For businesses seeking workload consolidation

based on Linux, the evaluation of Linux-enabled storage solutions and consolidation is the next natural step in a decidedly beneficial business process. Storage consolidation can also be improved using SAN solutions. We overview all IBM TotalStorage solutions during this chapter.

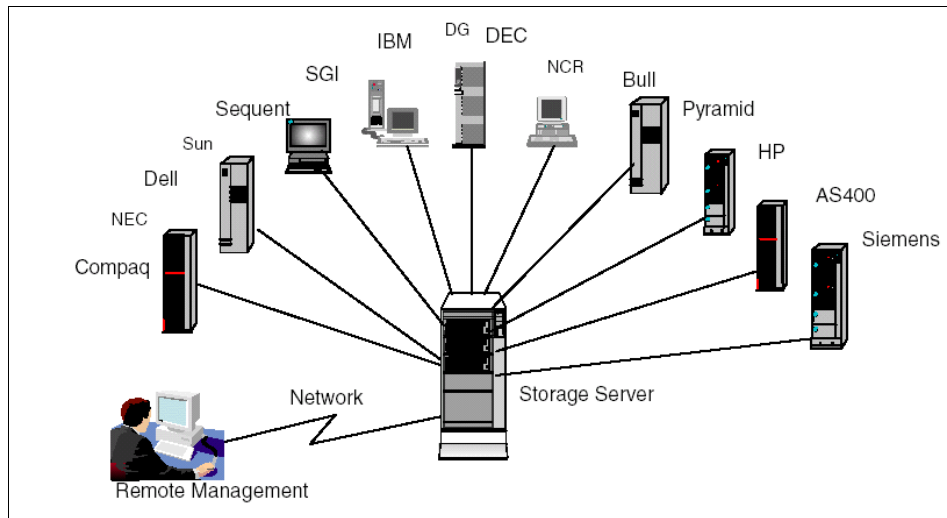


Figure 8-2 Physically centralized storage for virtually all servers

Linux delivers the power of open platform storage

For many organizations, one of the most appealing aspects of Linux is its open platform structure. Linux enables businesses to build capable, high-end systems using commercial, off-the-shelf and open-source software with easy attachments to a wide array of storage technologies, so companies can avoid investing in proprietary infrastructures that lock them in to a single or limited number of vendors.

Because of its open platform structure, Linux-enabled storage allows customers to expand their storage infrastructures to accommodate changing business needs within budget constraints — without paying a premium for proprietary operating systems with complicated licensing agreements. In this manner, Linux-enabled storage allows for predictable technology investment and low TCO.

Linux also delivers solid reliability, availability, and security. By combining the reliability legacy of UNIX systems with open source technology, Linux delivers a flexible platform to support even the most demanding business processes.

IBM lays the foundation for Linux-enabled storage

IBM delivers networked storage providing cost efficiencies and continuous operations to its customers, making IBM storage and Linux a natural fit. IBM TotalStorage products offer companies an extensive array of Linux-enabled storage hardware and software solutions, which in combination with the IBM product line meet virtually any customer need while exploiting the benefits of open-source, heterogeneous environments made possible through Linux.

IBM has one of the most extensive portfolios of Linux-enabled storage offerings in the industry, including direct attached storage, network attached storage (NAS), Storage Area Networks (SANs), tapes and storage management software, which all are ready for rapid deployment in pretested, documented, and supported configurations. Because IBM's Linux-enabled storage is prepared to support the reliability and performance required by mission-critical workloads, it provides an excellent foundation for workload consolidation.

8.2 IBM TotalStorage solutions

Linux-enabled storage is a requirement for all servers running Linux, regardless of the hardware infrastructure they rely on. We can observe that as Linux moves from supporting infrastructure workloads such as Web serving, networking, file and print operations, into the mission critical applications at the core of the data center, more stringent requirements are placed on the storage infrastructure to deliver superior reliability, scalability and security, and performance. These are not new requirements for IBM's storage portfolio, which has comprehensive Linux support now with a commitment to make it even stronger with our new products and solutions.

IBM continues to support Linux across all its server platforms, and leads the industry in storage networking based on open industry standards for heterogeneous platforms.

IBM is the leader in delivering world-class disk and tape systems, storage management software, services, and integrated solutions, and incorporates the following:

- ▶ Highly scalable storage - Enables administrators to manage growth and quickly respond to changing business needs, with ease of deployment and configuration
- ▶ High availability and fault-tolerant storage - Provides continuous and reliable access to data using technologies like RAID and clustering
- ▶ Improved data management - Helps administrators better control the security and growth of their data

- ▶ Increased storage utilization - Allocates storage through a centrally managed pool of storage
- ▶ Reduced administrative costs - Manage additional storage without having to add staff and, with IBM storage management tools, the ability to perform quicker problem resolution
- ▶ Platform independence - Enables sharing of data and possibility of simplification of backup procedures
- ▶ Centrally managed - Enables administrators to quickly respond to changing business needs, with ease of deployment and configuration
- ▶ Reduced administrative costs - Manage additional storage without having to add staff, and with IBM storage management tools, the ability to perform quicker problem resolution

8.3 IBM TotalStorage product line

This section provides a brief overview and positioning of IBM TotalStorage™ disk products that can be used and managed in a Linux environment focusing on this main product lines:

- ▶ IBM TotalStorage Enterprise Storage Server
- ▶ IBM TotalStorage FAStT Storage Server
- ▶ IBM TotalStorage tape drives
- ▶ IBM TotalStorage storage networking solutions

IBM's TotalStorage solutions supported on Linux are divided in three main lines: disk products, tape products, and storage networking solutions.

8.3.1 IBM TotalStorage disk products

IBM Totalstorage disk product have as common point to be based on hard disk technology. This product family has the following main members:

- ▶ **IBM TotalStorage Enterprise Storage Server (ESS):** This storage solution sets new standards in performance, automation and integration as well as capabilities that support continuous availability to data for the on demand world. It also supports many advanced functions, which can be critical for increasing data availability during planned outages and for protecting data from planned and unplanned outages. These advanced functions can provide important disaster recovery and backup protection.



Figure 8-3 TotalStorage disk based product family

- ▶ **FAStT 200 Storage Server:** Designed for workgroup and departmental servers that require an external storage solution. The single controller model provides a cost-effective solution, while the FAStT200 High Availability (HA) model features a fully redundant configuration with dual-active controllers. As your storage requirements grow, you can easily expand storage capacity by adding IBM FAStT EXP500 or EXP700 Expansion Units.
- ▶ **FAStT 600 Storage Server:** A mid-level storage server that can scale to over eight terabytes of fibre channel disk using 3 EXP700s, over sixteen terabytes of fibre channel disk with the Turbo feature using 7 EXP100s. It uses the latest in storage networking technology to provide an end-to-end 2 Gbps Fibre Channel solution.
- ▶ **FAStT 700 Storage Server:** Delivers superior performance with 2 Gbps Fibre Channel technology. The FAStT700 is designed to offer investment protection with advanced functions and flexible features. Scales from 36GB to over 32TB to support growing storage requirements created by e-business applications and offers advanced replication services to support business continuance.
- ▶ **FAStT 900 Storage Server:** Delivers breakthrough disk performance and outstanding reliability for demanding applications in compute-intensive environments. The FAStT900 is designed to offer investment protection with advanced functions and flexible features. Designed for today's on demand business needs, the FAStT900 easily scales from 36 GB to over 32 TB to support growing storage requirements. FAStT900 offers advanced replication services to support business continuance and disaster recovery. The

FAStT900 is an effective storage server for any enterprise seeking performance without borders.

8.3.2 IBM TotalStorage tape products

IBM TotalStorage tape products are based on the tape technology. The main tape products supported on Linux are:

- ▶ **LTO Ultrium Tape Drive 3580:** Offers breakthroughs in reliability, capacity and performance to meet the backup and archival needs for small to midrange environments. With its higher performance and capacity, the Ultrium Tape solution is an excellent alternative to S-DLT, DLT, 1/4 inch, 4 mm or 8 mm tapes drives.
- ▶ **LTO Ultrium Tape Autoloader 3581:** An ideal automation solution for handling the storage needs of small to medium sized environments. With its single Ultrium drive and up to seven tape cartridges, the Ultrium Tape Autoloader is designed to leverage the LTO technology to cost-effectively handle growing storage requirements.
- ▶ **Enterprise Tape Drive 3590:** Designed for automated tape libraries, including the IBM TotalStorage Enterprise Automated Tape Library 3494, IBM TotalStorage Virtual Tape Server, and StorageTek Silos. A 3494 with the 3590 E1A drive capacity is up to 1,122TB.

On library support products, IBM supports the following storage solutions:

- ▶ **Enterprise Tape Library 3494:** An excellent solution for today's large storage requirements. Modular, flexible, and reliable, the Enterprise Tape Library has a low entry price, is cost effective, and is backed by the service expertise of the IBM service organization. The Enterprise Tape Library offers a variety of models and features to fit all your needs and will grow as your business grows.
- ▶ **LTO Ultrium Scalable Tape Library 3583:** An ideal solution for cost-effectively handling a wide range of backup, archive, and disaster recovery data storage needs. The breakthrough reliability, capacity and performance of LTO offers an excellent alternative to DLT, 8mm, 4mm, or 1/4-inch tape drives for streaming data applications such as backup.
- ▶ **LTO UltraScalable Tape Library 3584:** An ideal solution for cost-effectively handling a wide range of backup, archive, and disaster recovery data storage needs. The breakthrough reliability, capacity and performance of LTO offers an excellent alternative to DLT, 8mm, 4mm, or 1/4-inch tape drives for streaming data applications such as backup.



Figure 8-4 IBM TotalStorage tape drive products

Note: find more information about IBM TotalStorage Linux-enabled products at: <http://www.storage.ibm.com/linux/index.html>

8.3.3 Comparing SAN X NAS technologies

Network storage combines the best features of networking and I/O—particularly the addressability, distance, and flexibility of networking and the raw performance and efficiency of I/O.

Network attached storage, or NAS, has a heritage rooted in ethernet-accessed data, and is mostly modeled after the network file server concept. NAS products typically provide integrated services in complete application packages that are compatible with the existing network infrastructure.

Storage Area Networks, or SANs, are new connectivity topologies designed to replace the existing SCSI I/O connection methods between systems and storage. SANs represent a new way of transporting data from the system that processes it to the place where it is stored

Table 8-1 SAN and NAS differences

	SAN	NAS
Protocol	Serial SCSI, FCP	CIFS, NFS
Network	Fibre Channel	Ethernet

	SAN	NAS
Source/Target	Server/Device	Client/Server, Server/Server
Transfer Objects	Device Blocks	Files
Storage Device Connection	Direct on Network	I/O BUS or Channel on Server
Embedded File System	No	Yes

8.3.4 IBM TotalStorage NAS products

Network attached storage products provide shared data to clients and servers over a LAN. Because of this, they are operating system independent. These products will work with Linux and with other OS.

IBM TotalStorage NAS products family

The IBM TotalStorage™ Network Attached Storage family is part of the overall IBM Storage Networking offering of hardware, software and services. IBM network attached storage (NAS) products are designed to enable users to increase the flexibility, efficiency and effectiveness of storage networking solutions and deployments. Based on an open-system concept, IBM NAS products are also designed to make future storage expansion as easy as installation is today.

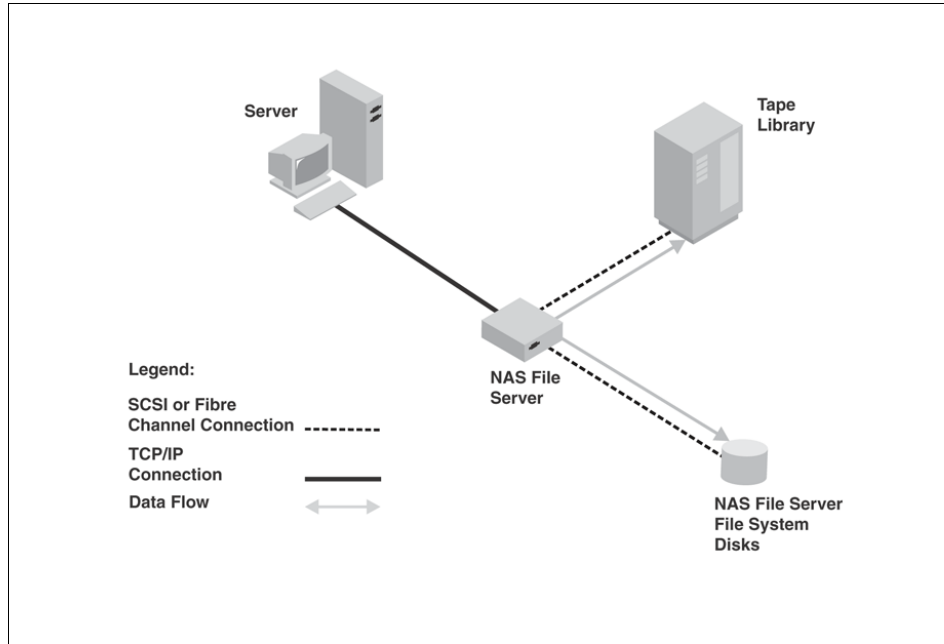


Figure 8-5 NAS sample configuration

The current TotalStorage NAS family includes the following products:

- ▶ **IBM TotalStorage NAS 100:** One of the first additions to the IBM TotalStorage™ Network Attached Storage (NAS) family, the IBM NAS 100 R12 features a 1U tabletop deployment or rack-mounting option. With a comprehensive systems management offering, it centrally manages storage for your distributed branch and remote locations. The IBM NAS 100 is designed to deliver an affordable, dedicated storage solution within a distributed or remote environment to support file sharing for office automation tasks, e-mail storage, file archiving or backup for connected clients.
- ▶ **IBM TotalStorage NAS 200:** The IBM TotalStorage™ Network Attached Storage (NAS) 200 Model 25T is one element in the overall IBM Storage Networking offering of hardware, software and services. The NAS 200 provides a powerful building block to increase the flexibility, efficiency and effectiveness of storage networking deployments. Additionally, because these NAS products are designed for interoperability, they can typically be integrated into existing IT infrastructures. IBM NAS products also scale, accommodating the quickly growing storage needs and changing technology demands of today's enterprises, which can ease the administrative burden on trained IT personnel.

- ▶ **IBM TotalStorage NAS 300:** The IBM TotalStorage™ Network Attached Storage 300 is part of the overall IBM Storage Networking offering of hardware, software and services. IBM network attached storage (NAS) products are designed to enable users to increase the flexibility, efficiency and effectiveness of storage networking solutions and deployments. Based on an open-system concept, IBM NAS products are also designed to make future storage expansion as easy as installation is today.
- ▶ **IBM TotalStorage NAS 300 Gateway:** The new IBM TotalStorage™ Network Attached Storage 300G series is part of the overall IBM Storage Networking offering of hardware, software and services. IBM network attached storage (NAS) products provide you with additional building blocks to increase the flexibility, efficiency and effectiveness of your storage networking solutions.




Figure 8-6 IBM Network Attached Storage family

We will overview the main features common on these products, and they are quite similar

Fast access to data

The IBM NAS servers are a tightly integrated collection of powerful subsystems assembled into a single appliance. Designed to provide NAS support for the departmental or small enterprise environment, the IBM NAS solutions are preloaded with optimized software and can operate 7 days a week, 24 hours a day. Redundant engines, dual RAID cards and dual Fibre Channel hubs are part

of a high-performance, faulttolerant design that helps minimize downtime caused by failure of any single element within the storage system.

Cross-platform data sharing

IBM NAS solutions are designed for use in both enterprise and departmental environments that have file-serving requirements across Windows or heterogeneous Linux, Windows and UNIX clients. Because of its multiple protocol support of CIFS (Windows), NFS (UNIX), HTTP, FTP, AppleTalk and NetWare, the IBM NAS products eliminate the need for a separate server for each protocol and enables clients to share files with protocols other than their own.

The IBM NAS solutions can help you perform the following mission-critical tasks:

- ▶ Remote or centralized order entry
- ▶ Payroll file storage and processing
- ▶ Accounts receivable storage and processing
- ▶ Communication network monitoring
- ▶ Customer support
- ▶ Backups of critical user data

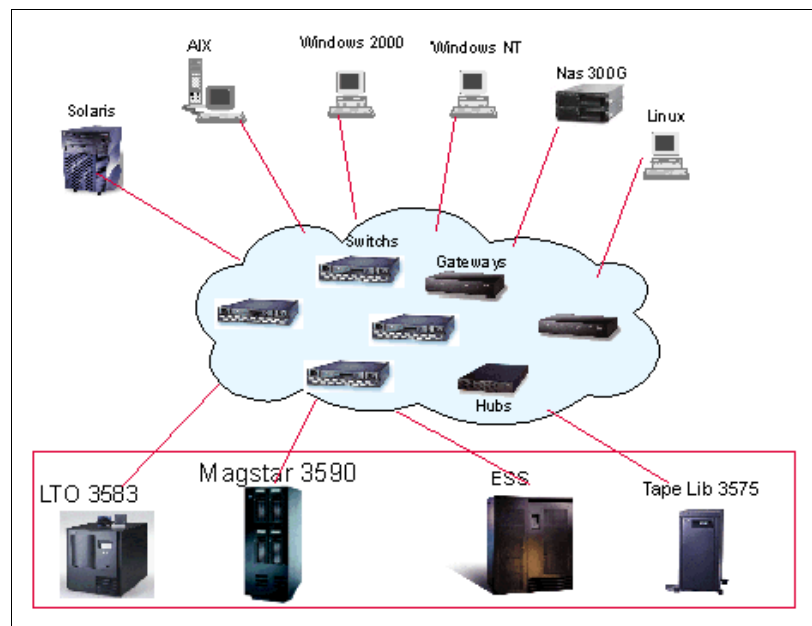


Figure 8-7 NAS implementation example

Designed for nondisruptive scalable growth

As storage networks become more complex, organizations must find cost-effective ways to manage the exponential growth of their storage needs. To

prepare for short and long-term storage requirements, organizations must have storage systems that can grow as their networking capacity needs increase. The IBM NAS products can easily scale to meet both current and future storage needs. Using Fibre Channel technology in its hubs, RAID controllers and hard disk drives, they are designed for performance, even in its maximum configuration. with other systems and to enhance performance. access to data. of ownership for either a small-enterprise or departmental environment. and UNIX clients. Channel hubs and engines help protect against system downtime.

IBM TotalStorage NAS product services

IBM Global Services provides the storage networking skills, and experience required for customers to take advantage of the latest storage networking technology to create a highly available storage environment that is easier to manage.

Product highlights

Available service offerings include planning, design and implementation services to ensure efficient transition to a reliable, highly available configuration for your business critical data.

IBM Global Services brings its breadth of expertise to help customers assess, plan and implement the best solution to the given business problem. Using trusted methods based on extensive industry experience, IBM Global Services can help customers quickly realize the maximum benefits of the IBM's Storage Networking products by ensuring the solution is integrated into the existing storage and management infrastructure.

Note: Find more about IBM services for TotalStorage NAS solutions at:

http://www.storage.ibm.com/snetwork/nas/pdf/networking_services_for_storage.pdf

http://www.storage.ibm.com/snetwork/nas/pdf/breadth_of_IGS_Storage_Networking_Services.pdf

8.3.5 IBM TotalStorage SAN products

Some of the products we have previously discussed can also be used as part of a Storage Area Network. On the server side, IBM offers a complete brand of SAN servers, for all type of business needs, covered with the complete IBM eServer brand:

- ▶ **zSeries:** At the heart of e-business infrastructure. Maximum performance, real-time responsiveness, application flexibility, and simplified management for data and transaction serving.
- ▶ **pSeries:** Combine the benefits of UNIX computing with IBM's leading-edge RISC technology in a broad product line from powerful workstations ideal for mechanical design up to parallel systems that can handle demanding scientific and technical computing, business intelligence and Web serving tasks.
- ▶ **xSeries:** Intel processor-based servers provide a range of configurations to meet your Windows, Linux, Novell and other system needs. It features IBM X-Architecture™, which is a design blueprint that leverages existing, innovative IBM technologies to build powerful, scalable and reliable Intel processor-based servers whether you support ten or tens of thousands of users
- ▶ **iSeries:** Run the most sophisticated business applications effortlessly and continuously. Cut complexity, slash bottom-line costs and speed implementation.

On the storage side, we can use those solutions, discussed previously on this chapter, and the ESS and FASTT more in detail later on this chapter:

- ▶ TotalStorage Enterprise Storage Server
- ▶ TotalStorage Enterprise Tape Library
- ▶ TotalStorage Linear Tape-Open
- ▶ TotalStorage Enterprise Tape System

IBM TotalStorage Virtualization family

The IBM TotalStorage Virtualization Family is designed to help you address the challenge to manage SAN-based storage. These solutions are designed to help simplify your storage infrastructure, optimize your storage utilization, and enable your business to adapt quickly and dynamically to variable environments. Virtualization is a key element to building an on demand operating environment.

On october 13th 2003, IBM announced the availability of the industry's first storage software technology designed to allow customers to easily share millions of data files in a heterogeneous environment and transform the management of data through automation. Based on the "Storage Tank™" technology developed by IBM Research, the IBM TotalStorage® SAN File System is designed to provide a single, centralized point of control to manage files and databases, which can help simplify administration and result in lower total costs.

Those IBM announcements on storage software solutions will be followed in 2004 by hardware products that will make use of this technology. We introduce this three main technologies:



Figure 8-8 IBM TotalStorage Virtualization Family

- ▶ **SAN File System:** The SAN File System marks a breakthrough in computing. The eagerly anticipated SAN File System is targeted at transforming the economics of SAN storage by allowing customers to better use existing hardware investments through a software virtualization layer.
- ▶ **SAN Volume Controller:** SVC is designed to reduce the complexity and costs of managing storage networks. It allows users to virtualize their storage and helps increase the utilization of existing capacity and centralize the management of multiple controllers in an open-system SAN environment. The SAN Volume Controller now supports attachment to non-IBM storage systems. Now storage administrators can reallocate and scale storage capacity and make changes to more underlying storage systems without disrupting applications.
- ▶ **SAN Integration Server:** The SAN Integration Server is designed to help integrate IBM virtualization technology, Fibre Channel switches and storage Redundant Array of Independent Disks (RAID) technologies into a preconfigured, comprehensive solution. Delivered and installed as a single unit, it offers upgrade options for connectivity, storage capacity, and performance levels. The solution was developed to provide the benefits of SAN with the ease of single-system manageability. SAN Integration Server will initially be capable of scaling to over 100 terabytes (TB) of storage capacity and connecting up to 42 hosts.

8.3.6 IBM TotalStorage Proven™ Program

IBM introduced the TotalStorage Proven program to help customers identify storage solutions and configurations that have been pre-tested for interoperability. This program is also intended to ease the burden on IBM Business Partners, independent software vendors (ISVs), and other storage solution providers with respect to independent interoperability testing.

The IBM TotalStorage Proven program builds on IBM's already extensive interoperability efforts to develop and deliver products and solutions that work together with third party products. Under the Storage Proven program, IBM will continue its work with hardware vendors, ISVs, and solution developers to test their products on IBM's extensive line of storage products.

The continuing benefits would include:

- ▶ Intended customer value:
 - Lower risk by investing in proven solutions
 - Lower Integration costs
 - Faster solution deployment

More information about storage proven can be found at:

<http://www.ibm.com/totalstorage/proven>

You can find an updated list of companies that have tested their products with IBM storage in a Linux environment at:

<http://www.storage.ibm.com/proven/solutions.htm>

8.3.7 ESS

This section introduces the IBM TotalStorage Enterprise Storage Server and discusses some of the benefits that can be achieved when using it. For more detailed information, please refer to the redbook: *IBM TotalStorage Enterprise Storage Server Model 800, SG24-6424*.

The IBM TotalStorage Enterprise Storage Server (ESS) is IBM's most powerful disk storage server, developed using IBM Seascope® architecture. The ESS provides unmatched functions for all the server family of e-business servers, and also for the non-IBM (that is, Intel-based and UNIX-based) families of servers. Across all of these environments, the ESS features unique capabilities that allow it to meet the most demanding requirements of performance, capacity, and data availability that the computing business may require. The Seascope architecture is the key to the development of IBM's storage products. Seascope

allows IBM to take the best of the technologies developed by the many IBM laboratories and integrate them, producing flexible and upgradeable storage solutions. This Seascape architecture design has allowed the IBM TotalStorage Enterprise Storage Server to evolve from the initial E models to the succeeding F models, and to the more recent 900 models, each featuring new, more powerful hardware and functional enhancements.

To meet the unique requirements of e-business, where massive swings in the demands placed on systems are common, and continuous operation is imperative, demands very high-performance, intelligent storage technologies and systems that can support any server application. The IBM TotalStorage Enterprise Storage Server has set new standards in function, performance, and scalability in these most challenging environments.




Figure 8-9 Front view of an open ESS 800

Figure 8-9 shows a photograph of an ESS Model 800 with the front covers removed. At the top of the frame are the disk drives, and immediately under them are the processor drawers that hold the cluster SMP processors. Just below the processor drawers are the I/O drawers that hold the SSA device adapters that connect to the SSA loops. Just below the I/O drawers are the host adapter bays

that hold the host adapters. At the bottom of the frame are the AC power supplies and batteries.

The ESS in this photo has two cages holding the disk drives (DDMs). If the capacity of this ESS was 64 or fewer disk drives, then the top right side of this ESS would have an empty cage in it. The photo clearly shows the two clusters, one on each side of the frame.

Between the two cages of DDMs is an operator panel that includes an emergency power switch, local power switch, power indicator lights, and message/error indicator lights.

For larger configurations, the ESS base enclosure attaches to an expansion enclosure rack that is the same size as the base ESS, and stands next to the ESS base frame.

ESS features and benefits

The ESS set a new standard for storage servers back in 1999 when it was first available, and since then it has evolved into the F models, and the recently announced third-generation ESS Model 800.

The IBM TotalStorage Enterprise Storage Server Model 800 introduces important changes that dramatically improve the overall value of ESS, and provides a strong base for strategic Storage Area Network (SAN) initiatives.

Storage consolidation

The ESS attachment versatility —and large capacity— enable the data from different platforms to be consolidated onto a single high-performance, high-availability box. Storage consolidation can be the first step towards server consolidation, reducing the number of boxes you have to manage, and allowing you to flexibly add or assign capacity when needed. The IBM TotalStorage Enterprise Storage Server supports all the major operating systems platforms, from the complete set of IBM server series of e-business servers and IBM NUMA-Q®, to the non-IBM Intel-based servers, and the different variations of UNIX based servers.

With a total capacity of more than 27 TB, and a diversified host attachment capability —SCSI, ESCON®, and Fibre Channel/FICON™— the IBM TotalStorage Enterprise Storage Server Model 800 provides outstanding performance while consolidating the storage demands of the heterogeneous set of server platforms that must be dealt with nowadays.

Performance

The IBM TotalStorage Enterprise Storage Server Model 800 integrates a new generation of hardware from top to bottom, allowing it to deliver unprecedented

levels of performance and throughput. Key features that characterize the performance enhancements of the ESS Model 800 are:

- ▶ New more powerful SSA device adapters
- ▶ Double CPI (Common Platform Interconnect) bandwidth
- ▶ Larger cache option (64 GB)
- ▶ Larger NVS (2 GB non-volatile storage) and bandwidth
- ▶ New, more powerful SMP dual active controller processors, with a Turbo feature option
- ▶ 2 Gb Fibre Channel/FICON server connectivity, doubling the bandwidth and instantaneous data rate of previous host adapters

Efficient cache management and powerful back-end

The ESS is designed to provide the highest performance for different types of workloads, even when mixing dissimilar workload demands. For example, zSeries servers and open systems put very different workload demands on the storage subsystem. A server like the zSeries typically has an I/O profile that is very cache-friendly, and takes advantage of the cache efficiency. On the other hand, an open system server does an I/O that can be very cache-unfriendly, because most of the hits are solved in the host server buffers. For the zSeries type of workload, the ESS has the option of a large cache (up to 64 GB) and —most important — it has efficient cache algorithms. For the cache unfriendly workloads, the ESS has a powerful back-end, with the SSA high performance disk adapters providing high I/O parallelism and throughput for the ever-evolving high-performance hard disk drives.

Data protection and remote copy functions

Many design characteristics and advanced functions of the IBM TotalStorage Enterprise Storage Server Model 800 contribute to protect the data in an effective manner.

Fault-tolerant design

The IBM TotalStorage Enterprise Storage Server is designed with no single point of failure. It is a fault-tolerant storage subsystem, which can be maintained and upgraded concurrently with user operation.

RAID 5 or RAID 10 data protection

With the IBM TotalStorage Enterprise Storage Server Model 800, there now is the additional option of configuring the disk arrays in a RAID 10 disposition (mirroring plus striping) in addition to the RAID 5 arrangement, which gives more flexibility when selecting the redundancy technique for protecting the users' data.

Storage Area Network (SAN)

The SAN strategy is to connect any server to any storage. As SANs migrate to 2 Gb technology, storage subsystems must exploit this more powerful bandwidth. Keeping pace with the evolution of SAN technology, the IBM TotalStorage Enterprise Storage Server Model 800 is introducing new 2 Gb Fibre Channel/FICON host adapters for native server connectivity and SAN integration.

These new 2 Gb Fibre Channel/FICON host adapters, which double the bandwidth and instantaneous data rate of the previous adapters available with the F Model, have one port with an LC connector for full-duplex data transfer over long-wave or short-wave fiber links. These adapters support the SCSI-FCP (Fibre Channel Protocol) and the FICON upper-level protocols.




Figure 8-10 Example of LAN-free data movement on a SAN configuration

The extensive connectivity capabilities make the ESS the unquestionable choice when planning the SAN solution. For the complete list of the ESS fabric support, please refer to:

<http://www.storage.ibm.com/disk/ess/supserver.htm>

For a description of the IBM TotalStorage SAN products, please refer to:

<http://www.storage.ibm.com/ibmsan/products/sanfabric.html>

ESS Specialist

The ESS Specialist interface (and the ESS Copy Services interface) consist of a set of Java applets that will provide you the interface to work with ESS volumes, partitions and all the configuration and setup procedures you may need in order to be able to work with an ESS. When you request a change to the configuration, the Java applets communicate with the microcode running on the ESS clusters to retrieve the current configuration data, submit the requested configuration change, and display the outcome of the request. This Web based interface make the ESS solution easy to use, and manageable remotely just from a Web browser.

ESS Copy Services and Linux

The Enterprise Storage Server Copy Services are enterprise-level functions that give you leading edge solutions to meet your needs for disaster recovery, data migration, and data duplication. The set of products which make up the ESS Copy Services suite are:

- ▶ **Peer-to-Peer Remote Copy (synchronous PPRC):** Peer-to-Peer Remote Copy (PPRC) is a synchronous remote data mirroring technology used primarily as part of a business continuance solution for protection against disk subsystem loss, or complete site failure. It is also used for remote migration of data and application workloads, and for off-site backups. It is a real-time copy technique that synchronously mirrors a primary set of volumes (that are being used by applications) onto a secondary set of volumes. Typically, the secondary volumes will be on a different ESS located at a remote location some distance away from the application site. Mirroring is done at a logical volume level.
- ▶ **Peer-to-Peer Remote Copy Extended Distance (non-synchronous PPRC-XD):** Peer-to-Peer Remote Copy Extended Distance (PPRC-XD) is a non-synchronous long-distance copy option for both open systems and zSeries servers. PPRC-XD can operate at distances well beyond the 103 km maximum supported distance for synchronous PPRC. It is an excellent solution for: data copy, data migration, off-site backup, transmission of data base logs and application recovery solutions based on periodic point-in-time copies of data
- ▶ **FlashCopy:** FlashCopy provides an instantaneous point-in-time copy of data, also called a T0 copy. FlashCopy enables applications that use either the source copy or the target copy to have them available almost immediately.

Only a minimal interruption to the application is required for the FlashCopy relationship to be established. The copy is then created under the covers by the IBM TotalStorage Enterprise Storage Server with minimal impact on other ESS activities. FlashCopy may be used in conjunction with either local or remote copies of data create by PPRC.

- ▶ **Extended Remote Copy (XRC):** XRC is a storage-based disaster recovery and workload migration solution that provides the capability to copy data in real time to a remote location. XRC is an extended function on all newer IBM storage controls, and operates in conjunction with the appropriate levels of DFSMS/MVS®. Since it is a DFSMS/MVS feature, it is not available on Linux.
- ▶ **Concurrent Copy (CC):** Concurrent Copy is a copy function that helps you keep your high data availability objectives by allowing point-in-time (PIT) copies of your data concurrent with normal application processing. Concurrent Copy works with the IBM TotalStorage Enterprise Storage Server and the DFSMS System Data Mover (SDM). Concurrent Copy is available for the z/OS and OS/390 operating systems, and requires software support provided in DFSMS/MVS®. This feature is not available on Linux because it requires DFSMS/MVS to work.

Note that XRC and CC are currently not available on Linux. These features may be available for Linux in a near future.

Peer-to-Peer Remote Copy (PPRC)

The Peer-to-Peer Remote Copy (PPRC) function is a hardware-based solution for mirroring logical volumes from a primary site (the application site) onto the volumes of a secondary site (the recovery site). PPRC is a remote copy solution for the open systems servers and for the zSeries servers.

Two modes of PPRC are available with the IBM TotalStorage Enterprise Storage Server Model 800:

- ▶ PPRC synchronous mode, for real-time mirroring between ESSs located up to 103 km apart
- ▶ PPRC Extended Distance (PPRC-XD) mode, for non-synchronous data copy over continental distances

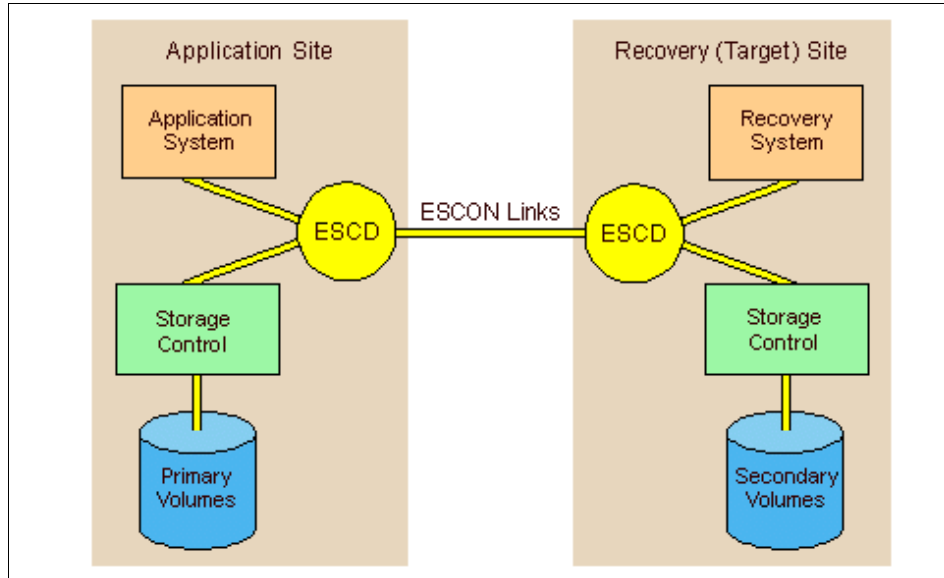


Figure 8-11 Synchronous volume copy PPRC

Extended Remote Copy (XRC)

Extended Remote Copy (XRC) is a combined hardware and software remote copy solution for the z/OS™ and OS/390® environments. The asynchronous characteristics of XRC make it suitable for continental distance implementations.

Point-in-Time Copy function

Users still need to take backups to protect data from logical errors and disasters. For all environments, taking backups of user data traditionally takes a considerable amount of time. Usually, backups are taken outside prime shift because of their duration and the consequent impact to normal operations.

Databases must be closed to create consistency and data integrity, and online systems are normally shut down.

With the IBM TotalStorage Enterprise Storage Server Model 800, the backup time has been reduced to a minimal amount of time when using the FlashCopy® function. FlashCopy creates an instant point-in-time copy of data, and makes it possible to access both the source and target copies immediately, thus allowing the applications to resume with minimal disruption.

8.3.8 FAStT

IBM Fibre Array Storage Technology (FAStT) solutions are designed to support the large and growing data storage requirements of business-critical applications.

The FAStT Storage Server is a RAID controller device that contains Fibre Channel (FC) interfaces to connect the host systems and the disk drive enclosures. The Storage Server provides high system availability through the use of hot-swappable and redundant components. The Storage Server features two RAID controller units, redundant power supplies, and fans. All these components are hot-swappable, which assures excellent system availability. A fan or power supply failure will not cause downtime, and such faults can be fixed while the system remains operational. The same is true for a disk failure if fault-tolerant RAID levels are used. With two RAID controller units and proper cabling, a RAID controller or path failure will not cause loss of access to data. The disk enclosures can be connected in a fully redundant manner, which provides a very high level of availability. On the host side FC connections, you can use up to four mini-hubs.

The FAStT Storage Server can support high-end configurations with massive storage capacities (up to 33 Tb per FAStT controller) and a large number of heterogeneous host systems. It offers a high level of availability, performance, and expandability.

FAStT models

At the time of writing this book, the FAStT900 and FAStT600 represented the newest addition to the FAStT family, complementing and joining the range of the earlier FAStT200, FAStT500, and FAStT700.

- ▶ The FAStT200 and FAStT600 are ideally placed for low-end mid-range setups such as workgroup and small department consolidations.
- ▶ The FAStT500 and FAStT700 are considered mid-range and are aimed at department storage consolidations, server clusters, and small to mid-range SANs.
- ▶ The FAStT900 is considered an entry-level enterprise product positioned for enterprise storage consolidation, mission critical databases, large SAN's, and high performance I/O.

IBM Storage Manager

To manage the FAStT storage server, use the IBM FAStT Storage Manager software. Refer to the RedBook 'Implementing Linux with IBM disk storage' on chapter 9., "FAStT Storage Manager" for more details.

<http://www.redbooks.ibm.com/redbooks/pdfs/sg246261.pdf>

At a glance, this software allows you to:

- ▶ Configure arrays and logical drives
- ▶ Assign your logical drives into storage partitions
- ▶ Replace and rebuild failed disk drives
- ▶ Expand the size of arrays and logical volumes
- ▶ Convert from one RAID level to another
- ▶ Perform troubleshooting and management tasks, like checking the status of FASTT Storage Server components, update the firmware or RAID controllers, and similar actions
- ▶ Configure and manage FlashCopy Volumes and Remote Volume Mirroring (FlashCopy and RVM are premium features that must be purchased)

It is also possible to use the serial interface and a terminal emulation utility. However, this is only meant for advanced troubleshooting and management. It should only be used when other management methods fail, and must be done under the supervision of IBM level 2 support.

FASTT Premium Features

IBM TotalStorage FASTT family supports Premium Features, those advanced features that must be purchased separately. Premium Features are a function of the IBM TotalStorage FASTT Storage Server firmware, and are usable from Linux attached hosts in the same manner as they would be from a host that is running any of the supported host operating systems. These additional features include:

- ▶ **Storage Partitioning:** Adds a high level of flexibility to the FASTT Storage Server. It allows you to connect a much higher number of host systems, either in standalone, or clustered mode.
- ▶ **FlashCopy:** A FlashCopy logical drive is a point-in-time image of a logical drive. It is the logical equivalent of a complete physical copy, but you create it much more quickly than a physical copy and it requires less disk space.
- ▶ **Remote Volume Mirroring:** The Remote Mirror option is used for online, real-time replication of data between storage subsystems over a remote distance. In the event of a disaster or unrecoverable error at one storage subsystem, the Remote Mirror option enables you to promote a second storage subsystem to take over responsibility for normal input/output (I/O) operations
- ▶ **VolumeCopy:** Used to copy data from arrays that use smaller capacity drives to arrays that use larger capacity drives, to back up data, or to restore FlashCopy logical drive data to the base logical drive.

Storage Partitioning

Storage Partitioning allows you to connect multiple host systems to the same storage server. It is a way of assigning logical drives to specific host systems or groups of hosts. This is known as *LUN masking*.

Logical drives in a storage partition are only visible and accessible by their assigned group or individual hosts. Heterogeneous host support means that the host systems can run different operating systems. However, be aware that all the host systems within a particular storage partition must run the same operating system because they will have unlimited access to all logical drives in this partition. Therefore, the file systems on these logical drives must be compatible with the host systems.



IBM Software for Linux

As part of the investment of IBM and strategy around the Linux operating system, IBM Software Group (SWG) is continually enlarging its software catalog for Linux. Under the DB2, Lotus, Tivoli, and WebSphere brands, IBM Software comprises the most comprehensive stack of middleware vital to successful deployments of information management, group collaboration, systems management, and Web services.

This chapter provides an overview of IBM Software on Linux through the following topics:

- ▶ An overview Linux and IBM Software availability
- ▶ WebSphere brand software for Linux
- ▶ DB2 Information Management software for Linux
- ▶ Tivoli brand software for Linux
- ▶ Lotus brand software for Linux
- ▶ Rational brand software for Linux
- ▶ Sample Linux solution deployments
- ▶ Customer references
- ▶ IBM Software for Linux resources

9.1 Overview of Linux and IBM Software availability

During the last few years, Linux has become the fastest growing general computing operating system in the market. It is now acknowledged to be an effective business solution platform.

Middleware is often the prominent driving force of business solutions, and is indispensable to an infrastructure. Middleware products are required to be available, flexible, and contain multiple functions. Furthermore, the market demands that they conform to open standards such as Java 2 Platform, Enterprise Edition (J2EE), and Web services. As a result, much of the IBM middleware catalog including DB2, WebSphere, Lotus, and Tivoli are now available for the Linux platform.

Business comes with risks. Therefore, it is important to select a software platform with proven reliability. It is also important to choose middleware that can execute well on many platforms including Linux. Both the operating system and applications that provide a solution to the customer's needs must be available now. Linux and IBM Software meet these criteria and offer a very powerful solution that includes:

- ▶ **Flexibility:** Linux allows you to adopt a variety of hardware and software that are appropriate to each business task. Linux can run on a variety of platforms along with IBM middleware, making it a good candidate for a multiplatform environment. Also, new applications are easily ported to an existing system.
- ▶ **Reliability:** The extensibility and availability of Linux and IBM Software have already been proven in customer environments. Combining IBM Software with the Linux operating system results in improved reliability, productivity, and customer satisfaction.
- ▶ **Cost performance:** IBM Software running Linux results in a reduced cost of the operating system, improved management of distributed systems, and a wide selection of compatible hardware. Also, running IBM Software on Linux greatly contributes to decreasing your total cost of ownership (TCO).

You can find the IBM Software for Linux home page on the Web at:

<http://www.ibm.com/software/os/linux/software/index.jsp>

Note: The *SWG Products, Solutions and Strategies Reference Guide*, G325-2132, provides an excellent overview of each of our software products. While this document is not specific to Linux-based products, it provides a wealth of information about each IBM software offering. To obtain a copy, go to the following Web site, select your country or region location, and search for the publication number:

<http://www.elink.ibm.link.ibm.com/public/applications/publications/cgi-bin/pbi.cgi>

IBM Software for Linux availability

IBM Software products for Linux are supported on a per architecture, per distribution basis. On the IBM Software for Linux Web site, you can find the most current list of IBM Software announcements. You can download a complete IBM Software products for Linux matrix on the Web at:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

Note: IBM Software that is listed as available for IBM Intel-based xSeries systems is also available for Intel-based systems from vendors other than IBM.

As a general statement of direction, IBM Software brands will support two Linux distributions: the Linux “enterprise” distributions from Red Hat (Red Hat Enterprise Linux) and the UnitedLinux consortium (SuSE Linux Enterprise Server (SLES), Turbolinux TLES, and Conectiva CLEE). Other distributions may be supported on an exception basis.

For a detailed description of each Linux distributor, see 1.9, “The Linux distributions” on page 24.

Software availability may also be affected by:

- ▶ Demand and target market for products on specific hardware architectures: For example, a market may demand a middleware product on IBM Intel-based xSeries servers but not on pSeries servers.
- ▶ Requirements on the Linux operating system by the software for performance or functionality: For example, many third-party software products require an “enterprise” version of Linux such as SLES rather than a “desktop” version of Linux such as the SuSE Linux Office Desktop.
- ▶ Market availability of IBM hardware platforms
- ▶ Market availability of Linux distributions

- ▶ Product release schedules of Linux distributors
- ▶ Internal IBM brand management decisions
- ▶ IBM middleware product introduction and obsolescence

Note: The absence of an IBM Software product in the software availability matrix is not an indication that the product will never be available for Linux. If you need an IBM Software product that is not listed in the software availability matrix, contact your local IBM sales representative.

9.2 On Demand business

An On Demand Business is an enterprise whose business process (integrated end-to-end across the company and with key partners, suppliers, and customers) can respond with speed to any customer demand, market opportunity, or external threat.

An On Demand Business is:

- ▶ Responsive
- ▶ Variable
- ▶ Focused
- ▶ Resilient

To meet these objectives, your operating environment must be On Demand too. In order to accomplish this business need, your On Demand Business environment is going to be integrated, automated, virtualized, and built on open standards. Implementing the operating environment as described is shown in Figure 9-1.

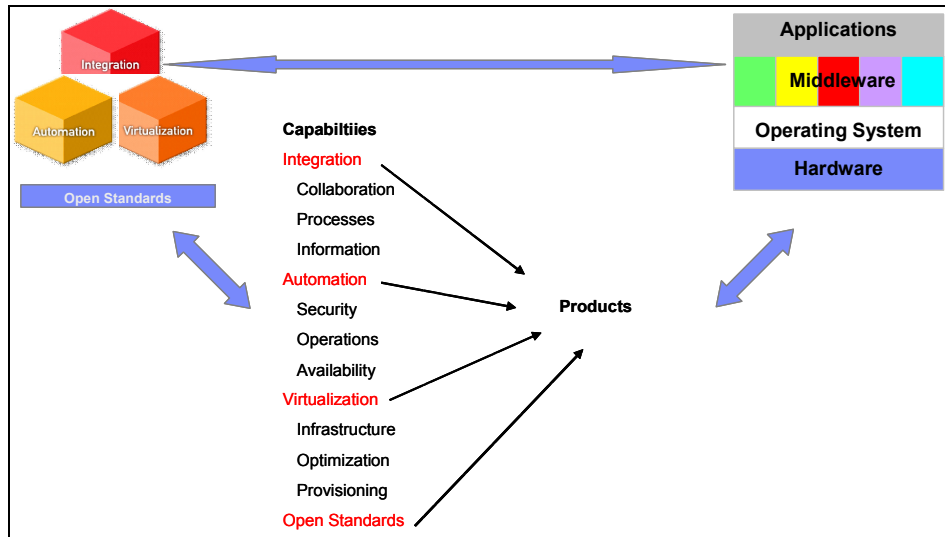


Figure 9-1 Implementing the operating environment

9.2.1 Open standards

Open standards deliver flexibility and avoid proprietary lock-in, offering you a choice of vendor. Open standards also enable integration between heterogeneous systems, both within the enterprise and with customers, partners, and suppliers.

Open standards and Linux

Linux provides a common API for applications and middleware across heterogeneous hardware, simplifying porting software to new hardware. It also provides a common operating system on a wide range of hardware platforms, from PDAs to mainframes.

Open standards and IBM Middleware

IBM Middleware uses open standards to provide a common environment for development, deployment, integration, and management of applications across heterogeneous platforms. For example, WebSphere Application Server provides a high-performance and extremely scalable transaction engine for dynamic e-business with support of Java 2 and Web services:

- ▶ J2EE 1.3
- ▶ Java SDK 1.4
- ▶ SOAP, UDDI, WSDL, XML

9.2.2 Integration

An integrated operating environment improves communication and collaboration within the enterprise and beyond, reacts quickly to changes in the marketplace by modifying business processes rapidly, and reacts in real-time to most relevant information by ensuring seamless flow of information.

Integration and Linux

Linux provides a common and open standard platform for application and middleware across heterogeneous hardware. It also can run alongside the existing operating system, integrating new applications with existing data and applications.

Integration and IBM Middleware

IBM Middleware provides integration of people, processes, and information across heterogeneous environments. The WebSphere Portal Server, for example, delivers a single point of personalized interaction with applications, content, processes, and people for a unified user experience.

9.2.3 Automation

Automation provides you secure access and control of information, resources, and applications; a scalable and consistent management of operations for end-to-end business systems; and avoids system failures by taking automated action to resolve problems.

Automation and Linux

Linux provides a reliable, robust, and secure operating system platform, and offers high availability through failover clustering. It also provides remote operations capability for distributed computing.

Automation and IBM Middleware

IBM Middleware manages security, availability, optimization, provisioning, and orchestration across heterogeneous environments. For example, IBM Tivoli Intelligent ThinkDynamic Orchestrator provides this functions within:

- ▶ Automating the provisioning and configuration of servers, operating systems, middleware, and applications
- ▶ Helping reduce costs, improve server utilization, and boost server-to-administrator ratio
- ▶ Anticipating plans, and dynamically providing server capacity to meet peak demands

9.2.4 Virtualization

Virtualization is a technique to consolidate and simplify infrastructure. It optimizes utilization and pool resources across heterogeneous environments. A virtualized environment is built in order to manage and provision system resources.

Virtualization and Linux

Linux offers server virtualization both multiple instances on one serve, and a single image of multiple servers through clusters and blades. It also enables sharing and optimization of pooled resources through grids and server consolidation.

Virtualization and IBM Middleware

IBM Middleware manages infrastructure simplification, storage optimization, grid enablement, and provisioning across heterogeneous environments. Solutions like DB2 Integrated Cluster Environment match this requisite providing a high-performance and reliable information management system that can scale from 2 to 1,000 nodes. It combines xSeries rack-optimized servers, Linux, and DB2 UDB in order to provide a virtualized operating environment while supporting open standards like XML.

9.3 IBM WebSphere brand products on Linux

IBM WebSphere® is the leading software platform for e-business on demand, delivering a proven, secure, and reliable software portfolio. Providing a responsive and comprehensive e-business platform, WebSphere software evolves to meet the demands of companies faced with challenging and continually changing business environments. Companies requiring a strong foundation and tools to reduce business risk, business portal software to strengthen business relationships, and business integration software to optimize operations. WebSphere platform products include WebSphere Application Server, WebSphere Commerce, WebSphere MQ, WebSphere Studio Application Developer. Figure 9-2 shows the three categories into which the WebSphere software platform can be divided.

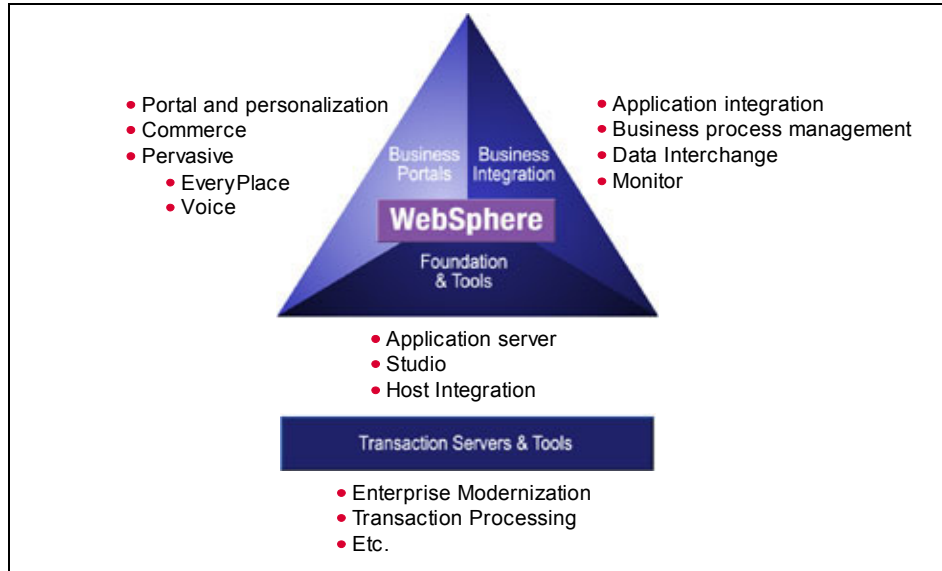


Figure 9-2 WebSphere platform overview

Each category is further explained here:

► **Business portals:**

This simplifies the process of creating a single interactive user experience. It allows you to deliver user-centric interactions for customers, partners, and employees alike through portal interfaces across all of your business processes. Commerce offerings enable full transactional support, and integrate with existing business systems. Expanding these user interactions for mobile access, providing a common experience throughout, delivers real value to e-business in our global and mobile marketplace.

► **Business integration**

This provides the ability to integrate data, applications, processes, and people. It gives businesses the power to collaborate with supply chain partners, leverage intelligence, and enhance customer relationship management. Leveraging existing capital investments is another benefit, since you integrate your current applications instead of installing all new ones. WebSphere's leadership in business integration is extended through security, performance, and scalability enhancements.

► **Foundation and tools**

Success is creating a foundation with flexibility and scalability for future growth. Through this category, WebSphere delivers business agility and scalability by leveraging Web services and existing software assets.

Comprehensive J2EE development and deployment capabilities are all necessities when you build for integration.

The foundation infrastructure, WebSphere Application Server Version 5, provides a rich e-business application deployment environment with a set of application services. These services include enhanced capabilities for transaction management in a heterogeneous environment, comprehensive Web services support, increased security, performance, availability, connectivity, and scalability. In addition, IBM WebSphere Studio provides an open, integrated application development environment that cuts development costs through unprecedented gains in productivity and quality.

► **Transaction servers and tools**

Success is achieved by capitalizing on existing environments. As new technology revolutionizes the way people work and interact, companies are left wondering what to do with their traditional software assets, applications, and data that run their business today.

The transactional environment provided by the WebSphere software suite allows you to integrate traditional core assets into a new technology infrastructure. It updates existing systems and leverages applications by transforming them into e-business components that can result in a new integrated e-business solution.

You can learn more on the WebSphere for Linux home page at:

<http://www.ibm.com/software/os/linux/software/webasp.shtml>

The value of WebSphere on Linux

The WebSphere product set is:

- **Flexible:** WebSphere is based on open standards for Java and Web services. This includes J2EE, Extended Markup Language (XML), Universal Description, Discovery, and Integration (UDDI), Simple Object Access Protocol (SOAP), and Web Services Description Language (WSDL). It supports the complexity of interlinking enterprises.
- **Reliable:** WebSphere provides a transactional environment for running Java servlets and Enterprise JavaBeans (EJB) on top of Apache Web servers.
- **Cost effective:** WebSphere reduces application development costs by reusing code. It optimizes operating costs because it can run on Linux from PC servers to server clusters to mainframes.

For more information about the IBM WebSphere family, see:

<http://www.ibm.com/software/info1/websphere/index.jsp>

9.3.1 WebSphere product availability on Linux

The following products are among the list of WebSphere-branded products on Linux:

- ▶ WebSphere Application Server
- ▶ WebSphere Business Components Composer
- ▶ WebSphere Business Components Studio
- ▶ WebSphere Commerce Business and Professional Edition
- ▶ WebSphere Commerce Developer Edition
- ▶ WebSphere Edge Server
- ▶ WebSphere Everyplace Connection Manager
- ▶ WebSphere Extended Personalization
- ▶ WebSphere Host on Demand
- ▶ WebSphere MQ
- ▶ WebSphere MQ Everyplace
- ▶ WebSphere MQ Integrator
- ▶ WebSphere MQ Express
- ▶ WebSphere Performance Pack Cache Manager
- ▶ WebSphere Portal for Multiplatform
- ▶ WebSphere Portal Express for Multiplatform
- ▶ WebSphere Portal Express Plus for Multiplatform
- ▶ WebSphere Site Analyzer
- ▶ WebSphere Studio Application Developer
- ▶ WebSphere Studio Device
- ▶ WebSphere Studio Homepage
- ▶ WebSphere Studio Application Monitor for z/OS
- ▶ WebSphere Studio Site Developer
- ▶ WebSphere Telecom Application Server and Toolkit
- ▶ WebSphere Transcoding Publisher

It is important to note that all versions of all products may not be available on all hardware platforms, or for all versions and distributions of Linux. For the most current availability and detailed descriptions of IBM Software products for Linux, download the IBM Linux software product matrix at:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

9.3.2 WebSphere Portal for Multiplatforms

IBM WebSphere Portal allows people to interact with the on demand world in a personalized way. They can automatically get the dynamic information they need. They can quickly execute business processes across critical applications. They can collaborate with portal users inside and outside your e-business. By providing these industry-leading portal solutions for your e-business, IBM helps

to improve employee productivity, cuts costs, and strengthens relationships with customers and trading partners.

The WebSphere Portal for Multiplatforms offerings are modular in nature with each offering building upon a common architecture and capabilities. The IBM WebSphere® Portal Enable for Multiplatforms offering is the basic edition of WebSphere Portal for Multiplatforms. It helps you quickly build scalable portals to simplify and accelerate access to personalized information and applications. WebSphere® Portal Extend for Multiplatforms, Version 5.0 includes all the robust features of WebSphere Portal Enable, and introduces collaboration capabilities, enterprise search functions, and portal usage analysis. These functions help you improve employee productivity and continually strengthen relationships with your customers and trading partners.

For more information, visit:

<http://www.ibm.com/software/genservers/portal/features/>

ftp://ftp.software.ibm.com/software/genservers/portal/Enable_G325-2112-01.pdf

ftp://ftp.software.ibm.com/software/genservers/portal/Extend_G325-5463-03.pdf

Figure 9-3 shows a simple production scenario. It demonstrates how a complete Internet and intranet Web environment can improve e-business with your partners, customers, suppliers, or employees.

WebSphere Portal Content Publishing ships with IBM WebSphere Portal for Multiplatform. The software is made up of browser-based Web content management tools that allow your non-technical users to manage the content within WebSphere Portal.

With WebSphere Portal Content Publish you can:

- ▶ Manage workflow
- ▶ Develop content
- ▶ Publish content

WebSphere Portal for Multiplatform also provides a feature to manage documents across your company. WebSphere Portal Document Manager enhances work productivity, frees client machine space, and streamlines complicated document management processes. It provides a centralized location for documents, and built-in methods to track changes and comments from members of the work team.

With this feature, a Web-based word processor, spreadsheet, and a presentation are provided to use the portal as a document builder.

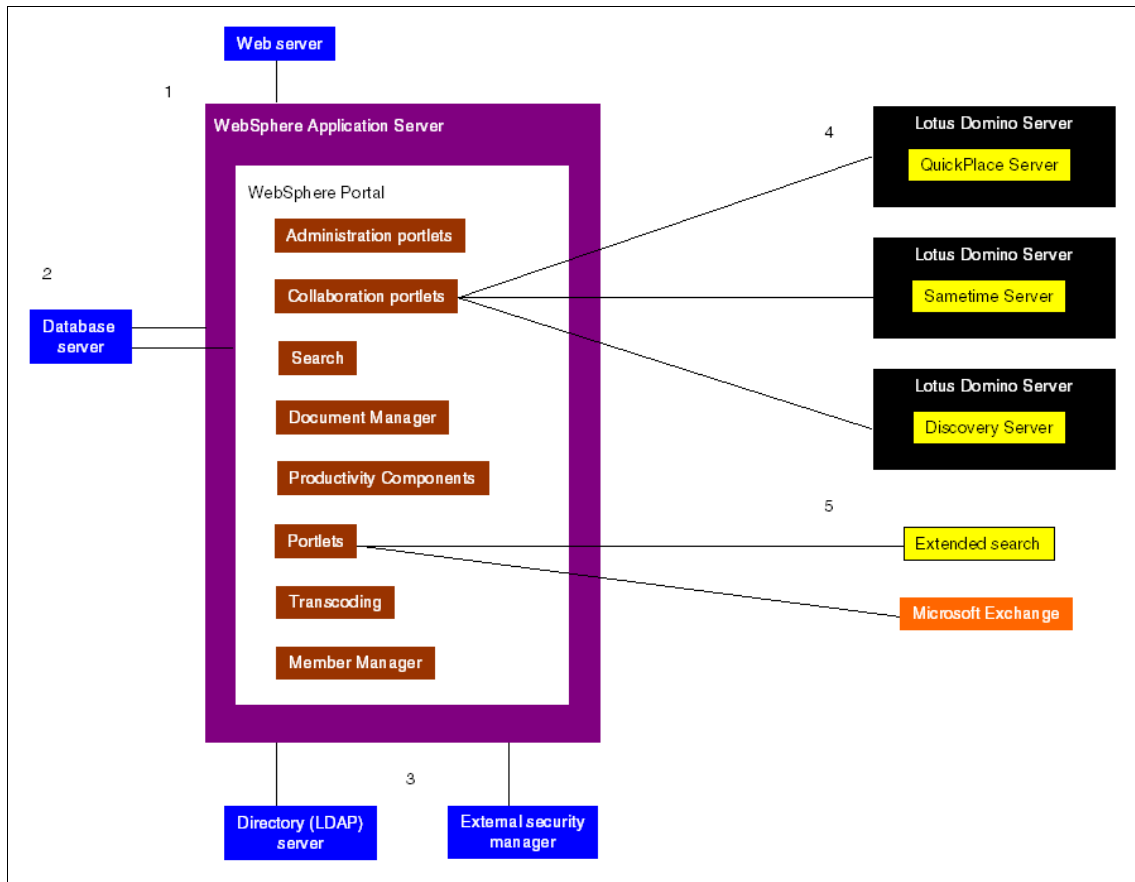


Figure 9-3 WebSphere Portal software topology

9.3.3 Linux-related WebSphere resources

This section provides information about where to find additional WebSphere resources.

Evaluation copy of WebSphere software

Free evaluation copies of certain Linux-enabled IBM middleware products, complete with installation, configuration, and support documentation are available through the Web. You can also obtain evaluation copies of WebSphere Advanced Server, WebSphere Studio Site Developer, WebSphere MQ, and Linux-enabled IBM Software.

To learn more about how to obtain your own evaluation copies, see:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

WebSphere Developer Domain

IBM provides a valuable set of online technical developer resources through the WebSphere Developer Domain Web site. You can find this Web site, which often includes new technology previews and WebSphere developer experiences, at:

<http://www.ibm.com/developerworks/websphere/>

Linux-specific items available from the WebSphere Developer Domain include:

- ▶ **ThreadAnalyzer technology preview:** ThreadAnalyzer gathers and analyzes thread dumps from WebSphere Application Server, Version 5. It can obtain a thread dump or open an existing thread dump. ThreadAnalyzer is supported on Solaris, AIX, Windows, and Linux on Intel and 390. You can download the ThreadAnalyzer technology preview from the Web at:
http://www.ibm.com/developerworks/websphere/downloads/thread_analyzer.html
- ▶ **Web services:** IBM WebSphere Web services for J2EE technology preview supports emerging Java Web services standards such as Java application programming interface (API) for XML-based RPC Version 1.0 (JAX-RPC) and Web services for J2EE. It contains an implementation of the emerging WS-Security standard, and is available for use with WebSphere Application Server Version 5.0. You can download the Web Services from the Web at:
http://www-106.ibm.com/developerworks/websphere/downloads/web_services.html
- ▶ **WebSphere Application Server Performance Advisor Technology Preview:** This package consists of two performance advisors to help tune systems for optimal performance. The advisors use the Performance Monitoring Infrastructure (PMI) data to suggest configuration changes to the WebSphere thread pools, connection pools, prepared statement cache, session cache, and heap size. You can download the Web services from:
http://www.ibm.com/developerworks/websphere/downloads/performance_advisors.html
- ▶ **Developing Enterprise Java Applications using DB2 Universal Database (UDB) Version 8:** This technical article summarizes enhancements in DB2 UDB Version 8 for Linux, UNIX, and Windows for Java programming. This includes integration with WebSphere, development tools, XML technology,

and information about the new JDBC Type 4 driver. You can download this article from the Web at:

<ftp://ftp.software.ibm.com/software/data/pubs/papers/db2aj.pdf>

- ▶ **Migrating to WebSphere Application Server Version 5:** This overview is designed to help you get started migrating your Web applications and configuration from earlier versions of WebSphere Application Server to Version 5. The tasks are described at a higher level to give you the big picture on some migrating essentials, including preparation guidelines, important migration options, and a first look at the Version 5 administrative server. You can find the complete document at:

http://www.ibm.com/developerworks/websphere/library/techarticles/0304_chalmers/chalmers.html

- ▶ **Highly available infrastructure with WebSphere Network Deployment:** This article discusses how to implement a highly availability infrastructure for the Deployment Manager component of WebSphere Application Server Network Deployment, Version 5. Learn how to do this without purchasing multiple servers and hardware clustering software for each component in the infrastructure. For the complete technical article, see:

http://www.ibm.com/developerworks/apps/transform.wss?URL=/developerworks/websphere/library/techarticles/0304_alcott/alcott.xml&xslURL=/developerworks/websphere/xsl/document.xsl&format=one-column

Relevant Linux-related IBM Redbooks

The current list of relevant WebSphere on Linux IBM Redbooks and Redpapers includes:

- ▶ *CICS Transaction Gateway V5 The WebSphere Connector for CICS*, SG24-6133
- ▶ *Linux Application Development Using WebSphere Studio 5*, SG24-6431
- ▶ *WebSphere Application Server Test Environment Guide*, SG24-6817
- ▶ *IBM WebSphere Portal V4.1 Handbook Volume 1*, SG24-6883
- ▶ *WebSphere Commerce V5.4 for Linux, Infrastructure and Deployment Patterns*, REDP0214
- ▶ *IBM WebSphere V5.0 for Linux, Implementation and Deployment Guide - WebSphere Handbook Series*, REDP3601
- ▶ *Linux on IBM zSeries and S/390: z/VM configuration for WebSphere Deployments*, REDP3661
- ▶ *WebSphere Solution Guide: WebSphere Application Server Express V.5*, SG24-6925

The list of Linux-related IBM Redbooks and Redpapers is continually being updated. You can find these documents on the Web at:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

9.4 IBM Information Management products on Linux

IBM has the richest set of information management products on the market. Many of the components of the three major information management brands such as DB2, Informix, and U2 are available on Linux.

9.4.1 DB2 product family

IBM DB2 database software is the worldwide market share leader in the industry. It marks the next stage in the evolution of the relational database. DB2 is the industry's first multimedia, Web-ready relational database management system (RDBMS) that delivers leading capabilities in reliability, performance, and scalability. It is the database of choice for customers and partners who are developing and deploying critical solutions.

You can find detailed product information for DB2 on Linux on the Web at:

<http://www.ibm.com/db2/linux>

The IBM DB2 product can add significant value to and benefit a customer business in the following ways:

- ▶ Lowest TCO
- ▶ Broad platform support
- ▶ Seamless connectivity to other data
- ▶ World wide 24-x-7 service and support
- ▶ IBM and independent software vendor (ISV) application and tool enablement
- ▶ Reliability, Scalability, Availability (RAS) and high performance
- ▶ DB2 self management and smart database technology to reduce the complexity

Specific value propositions for DB2 for zSeries

zSeries customers can extend and leverage their current investment in hardware. Hundreds or even thousands of Linux operating systems can coexist with other operating systems on a single zSeries under VM. With VM, a customer can offer a complete Linux server environment to each application developer, and still run host production systems on the same zSeries server.

You can deploy new applications on Linux and still connect to their existing DB2 data using DB2 Connect. You can also use IBM DB2 Connect on Linux for zSeries to consolidate all gateway connectivity processing at the host. This eliminates the need for middle-tier hardware.

Customers can take advantage of the features specific to UNIX environments on DB2 UDB on Linux. With DB2 UDB and DB2 Connect on Linux for zSeries, it is possible for vendor applications that are currently available for UNIX customers to leverage the zSeries platform.

Specific value propositions for DB2 for xSeries (Intel-Architecture server)

xSeries customers can reduce their TCO. DB2 UDB on Linux for xSeries can run on low-cost hardware. Also, customers are not locked into a proprietary operating system. DB2 Enterprise Server Edition with Data Partitioning Feature (ESE + DPF) for Linux allows customers to add extra capacity as needed for Web serving and data warehousing at a lower price than other clustering solutions that are currently available.

Moreover, DB2 is the performance leader on Linux. DB2 is among the gtp 2 in a recent industry benchmark, TPC-H, from 22 November 2002. For more information, see:

http://www.tpc.org/tpch/results/tpch_perf_results.asp

DB2 product availability on Linux

The list of DB2-branded products on Linux includes:

- ▶ DB2 UDB (Workgroup, Extended Edition, Enterprise Server Edition, Personal)
- ▶ DB2 Express (special offering for small and medium businesses (SMB))
- ▶ DB2 Connect
- ▶ DB2 Everyplace
- ▶ DB2 Intelligent Miner™ Scoring
- ▶ DB2 Net Search Extender
- ▶ DB2 Relational Connect
- ▶ DB2 Content Manager
- ▶ DB2 Cube Views
- ▶ DB2 Performance Expert
- ▶ DB2 Universal Database Data Warehouse

For the most current availability and detailed descriptions about IBM DB2 Software products for Linux, download the IBM Linux software product matrix at:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

IBM DB2 Integrated Cluster Environment (DB2 ICE)

The IBM DB2 Integrated Cluster Environment (DB2 ICE) for Linux is a completely integrated, high-performance, and pre-tested solution that incorporates best-of-breed software, hardware, and services. It provides you with a high performance, reliable data management system that can scale from 2 to 1,000 nodes. You can order a pre-configured solution or customize it to meet your needs.

The core solution is based on IBM DB2 Universal Database and IBM Linux Cluster 1350 (xSeries, 325, BladeCenter). The other components in a DB2 Integrated Cluster Environment can include (but not limited to):

- ▶ IBM WebSphere Application Server
- ▶ IBM Tivoli System Software
- ▶ IBM TotalStorage
- ▶ IBM Global Services
- ▶ SuSE or Red Hat Linux
- ▶ InfiniBand products from Mellanox and Voltaire

DB2 UDB support for Linux

DB2 UDB for Linux runs on a variety of hardware platforms. In addition to the Intel 32-bit hardware, DB2 for Linux today runs on the zSeries, pSeries, and iSeries. A Linux Validation Program available on the Web provides a process to ensure that DB2 UDB operates successfully within Linux environment. The goal of this program is to specifically identify those hardware and software products for Linux systems that work best with DB2. For more information about this program, see:

<http://www.ibm.com/software/data/db2/linux/validate/progdesc.html>

On the 32-bit x86 platform, DB2 requires systems equipped with Intel Pentium, Xeon, or AMD Athlon or later processors. DB2 Version 7.2 is supported on any Linux distribution with 2.4 kernel, glibc 2.1.2, and libstdc++ 2.9.0 or later. With DB2 Version 8, the Linux Validation Program is introduced to ensure better inter-operability of DB2 on Linux platforms and to enhance customer experience.

IBM is working closely with the Linux community, Linux Distribution Partners, and independent software and hardware vendors to test DB2 using various Linux kernels, distributions, third-party products, hardware, and other components that interact with DB2. As a part of the validation program, before a product is supported with DB2, it undergoes rigorous testing. This gives users confidence to run DB2 with other products right out of the box.

As an enhancement to the general Linux distribution support statement for IBM Software, DB2 UDB is supported for a large number of Linux distributions. The

Linux distributions that are successfully validated for use with DB2 UDB include Red Hat Linux Professional 7.2, 7.3, and Red Hat Enterprise Linux 3; SuSE Professional 7.3, 8.0, and 8.1; SuSE Linux Enterprise Server 7 and 8; Turbolinux 7 Server, Turbolinux 8 Server, Turbolinux Enterprise Server 8.0; and other distributions powered by UnitedLinux 1.0. See the following validation Web site for a complete and current table of the supported platforms and distributions:

<http://www.ibm.com/software/data/db2/linux/validate/>

Note: To run Java applications and the graphical tools that come with DB2 (such as Control Center), JDK 1.3.1 is required. For convenience, it is shipped on the DB2 Version 8.1 CDs. In most cases, it is automatically installed with DB2 installation. Refer to the validation Web site (above) for the most current list of supported products.

9.4.2 Informix product family

The Informix information management product set is widely deployed on various distributed servers such as UNIX and Windows. It provides the backbone of information management for many small to medium customers.

You can find detailed information about Informix products on the Web at:

<http://www.ibm.com/software/data/informix/>

Informix product availability on Linux

The Informix brand of products on Linux includes:

- ▶ Informix 4GL Compiler (plus Interactive Debugger)
- ▶ Informix 4GL Rapid Development System (plus Runtime)
- ▶ Informix C-ISAM
- ▶ Informix Client Software Development Kit (CSDK)
- ▶ Informix Connect Runtime
- ▶ Informix D4GL
- ▶ Informix Dynamic Server - IDS (Workgroup, Enterprise and Unlimited)
- ▶ Informix ESQL (Cobol, C)
- ▶ Informix Excalibur Text Search Datablade
- ▶ Informix Extended Parallel Server (XPS)
- ▶ Informix Geodetic Datablade
- ▶ Informix Image Foundation Datablade Module
- ▶ Informix Internet Foundation 2000 (IIF2K)
- ▶ Informix NET
- ▶ Informix OnLine
- ▶ Informix Red Brick™ Warehouse

- ▶ Informix SE Personal Edition
- ▶ Informix Server Administration
- ▶ Informix Spacial Datablade Module
- ▶ Informix TimeSeries Datablade
- ▶ Informix TimeSeries Realtime Loader
- ▶ Informix Web Datablade

For the most current availability and detailed descriptions of Informix software products for Linux, download the IBM Linux software product matrix at:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

Informix support for Linux

Linux is available from multiple distributors including Red Hat, SuSE, and Turbolinux. Informix moved from supporting each distributor's version, to supporting the Linux component versions. This allows the Linux community to define the distributors that support Informix products.

Regardless of the Linux distributor, a formal recertification of Informix product releases on the Linux operating system is only required when a major Linux upgrade is formally released. This is indicated by the changing of the second decimal point number of the kernel or glibc version.

For example, a certification from kernel version 2.2.5 to 2.2.6 is not required. However, a formal certification from kernel 2.2.5 to 2.4.0 is required. A certification from glibc 2.0.7 to 2.0.8 is not required, but it is required when the version changes from 2.0.7 to 2.1.0.

9.4.3 U2 product family

The U2 information management family products are similar to those found in a pick operating system. U2 offers high performance and scalable information management environments for embedding in vertical applications. Over 3.9 million users rely on IBM U2 Information Management solutions.

You can find detailed information about IBM U2 products on the Web at:

<http://www.ibm.com/software/data/u2/>

U2 product availability on Linux

The U2-branded products on Linux include:

- ▶ IBM Cobol Direct Connect
- ▶ MITS
- ▶ IBM RedBack®

- ▶ IBM SystemBuilder™ Plus
- ▶ IBM Unidata
- ▶ IBM Universe

For the most current availability and detailed descriptions of IBM Software products for Linux, download the IBM Linux software product matrix from:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

9.4.4 IBM Information Management resources

The following sections provide information about other resources that are available for information management.

DB2 UDB and DB2 Migration Kit: Evaluation copy

Free evaluation copies of certain Linux-enabled IBM middleware products, complete with installation, configuration, and support documentation are available through the Web. You can download evaluation copies of DB2 UDB and the DB2 Migration Kit, as well as other Linux-enabled IBM Software, from:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

IBM DB2 Migration Toolkit for Linux

The IBM DB2 Migration Toolkit (MTK) V1.1 helps you migrate from Sybase and Microsoft SQL Server (Versions 6 and 7) to DB2 UDB Version 7.2 databases on any supported DB2 UDB workstation platform. With the Migration Toolkit, you can automatically convert data types, tables, columns, views, indexes, stored procedures, and triggers into equivalent DB2 database objects.

The Migration Toolkit provides database administrators (DBAs) and application programmers with the tools needed to automate previously inefficient and costly migration tasks. You can reduce downtime, eliminate human error, and reduce person hours, and other resources associated with traditional database migration by using the features in the Migration Toolkit.

You can download the IBM DB2 Migration Toolkit from IBM at:

<http://www.ibm.com/software/data/db2/migration/mtk/>

The IBM DB2 Migration Toolkit Version 1.2 is available for download from IBM at:

https://www.software.ibm.com/dl/db2_03/migtkit-p

Tech Guide: DB2 for Linux

This IBM sponsored white paper from Newmediary, Inc. describes DB2 products on Linux, the market for databases on Linux, and comparisons with open source alternatives. You can download it free (after registering) from the Web at:

<http://www.ibm.com/developerworks/db2/library/techarticle/0202applied/0202applied.html>

Information Management Developer Works Domain

IBM provides a valuable set of online technical developer resources through Developer Works Domain Web sites. These Web sites, which often include new technology previews and developer experiences, are listed according to product family name.

You can find the DB2 Developer Works Domain on the Web at:

<http://www.ibm.com/developerworks/db2/>

And you can find the Informix Developer Works Zone on the Web at:

<http://www.ibm.com/developerworks/db2/zones/informix/>

DB2 UDB and the highly available data store

Critical database applications demand a robust strategy to prevent data loss and guarantee the high availability of your data store. This article surveys your options for high availability on Linux. You can find the complete article at:

<http://www.ibm.com/developerworks/db2/library/techarticle/0310melnyk/0310melnyk.html>

DB2 porting Web site

An IBM DB2 UDB porting site is now available on the Internet. This site is intended to help you find the information you need to port an application and its data from other database management systems to DB2 UDB. You can find the DB2 porting Web site at:

<http://www.ibm.com/developerworks/db2/zones/porting/>

Relevant Linux-related Redbooks

For more information about Linux, consult the following IBM Redbooks and Redpapers:

- ▶ *Linux Web Hosting with WebSphere - DB2 and Domino*, SG24-6007
- ▶ *DB2 Evaluation Guide for Linux*, SG24-6934

- ▶ *Linux for WebSphere and DB2 Servers*, SG24-5850
- ▶ *e-Business Intelligence: Leveraging DB2 for Linux on S/390*, SG24-5687
- ▶ *e-Business Intelligence: Data Mart Solutions with DB2 for Linux on zSeries*, SG24-6294
- ▶ *Up and Running with DB2 for Linux*, SG24-6899
- ▶ *DB2 Universal Database Enterprise Server Edition on a Linux Cluster System*, REDP-3640

For the latest Linux-related IBM Redbooks and Redpapers, see:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

You can also refer to the article *Web Click Stream Analysis using Linux Clusters*, which is located on the Web at:

<http://www-106.ibm.com/developerworks/db2/library/techarticle/lurie/0111lurie.html>

9.5 IBM Tivoli products on Linux

Tivoli offerings consist of four feature solution sets:

- ▶ **Performance and availability management:** IBM Tivoli performance and availability management solutions provide the central nervous system for your complicated e-business landscape. They constantly gather information about your hardware, software, and network devices, and in many cases will cure problems before they actually occur.
- ▶ **Configuration and operation management:** The IBM Tivoli configuration and operations solution delivers reliable services to your end users and makes them more productive. This suite of applications distributes software reliably and securely, manages the change and control of IT assets, automates workflow through the enterprise, and remotely controls systems and applications.
- ▶ **Security management:** IBM Tivoli security management solutions address two critical e-business challenges: automated identity management and security event management. The IBM Tivoli identity management solution helps you quickly realize return on investment (ROI) by bringing users, systems, and applications online fast, while effectively managing users, access rights, and privacy preferences throughout the identity life cycle. The IBM Tivoli security event management solution helps you actively monitor, correlate, and quickly respond to IT security incidents across your e-business.

- **Storage management:** IBM Tivoli Storage Management protects your organization's data from hardware failures and other errors by storing backup and archive copies of data on offline storage. Scaling to protect thousands of computers running a dozen operating system platforms, its intelligent data movement and store techniques and complete automation, reduce administration costs while increasing service levels. Tivoli storage products provide a combination of scalability, intelligent data technology, disaster preparation, and broad platform and application support through one centralized, automated solution.

Figure 9-4 illustrates the relationship of these solution sets.

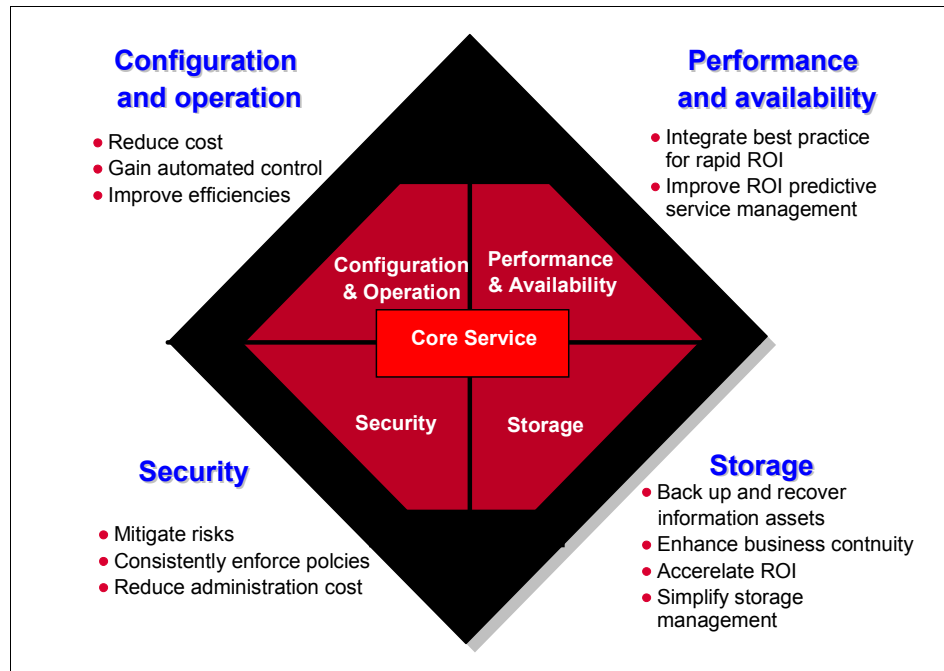


Figure 9-4 Tivoli IT management solutions

Table 9-1 shows to you a sample of offerings that are now available for Linux. For a complete lists of Tivoli products available for Linux, download the IBM Software Linux Matrix from:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

Table 9-1 Samples of offerings for Linux

Performance and availability	Configuration and operation	Security	Storage
Enterprise Console	Workload scheduler	Access Manager	Storage Manager
Monitoring	Configuration Manager	Identity Manager	Storage Area Network (SAN) Manager

IBM Tivoli Enterprise Console

Available on the zSeries and xSeries platforms, this product provides a centralized point of control that keep your IT staff in close and efficient control of events happening across all systems and networks. For more information, see:

<http://www.ibm.com/software/tivoli/products/enterprise-console/>

IBM Tivoli Monitoring

Available on the zSeries and xSeries platforms, this product provides monitoring for essential system resources to detect bottlenecks and potential problems, and to automatically recover for critical situation. For more information, see:

<http://www.ibm.com/software/tivoli/products/monitor/>

IBM Workload Scheduler

Available on the zSeries and xSeries platforms, this product automates, monitors, and controls the flow of work through your enterprise's entire IT infrastructure on both local and remote systems. For more information, see:

<http://www.ibm.com/software/tivoli/products/scheduler/>

IBM Tivoli Configuration Manager

Available on the zSeries and xSeries platforms, this product delivers an integrated solution for deploying software and for tracking hardware and software configurations. For more information, see:

<http://www.ibm.com/software/tivoli/products/config-mgr/>

IBM Tivoli Access Manager

Available on the zSeries platform, this product provides end-to-end security for e-business. This includes Web single signon (SSO), distributed Web-based administration, and policy-based security. For more information, see:

<http://www.ibm.com/software/tivoli/products/access-mgr-e-bus/>

IBM Tivoli Identity Manager

This product provides centrally coordinated creation of user accounts and automated workflow for approval. For more information, see:

<http://www.ibm.com/software/tivoli/products/identity-mgr/>

IBM Tivoli Storage Manager

This product provide automated data backup and restore functions. It also centralizes storage management operations. For more information, see:

<http://www.ibm.com/software/tivoli/products/storage-mgr/>

IBM Tivoli Storage Area Network Manager

This product is a comprehensive management solution for multivendor SANs. It includes automatic resource and topology discovery, monitoring alerts, zone control, and SAN error prediction capabilities. For more information, see:

<http://www.ibm.com/software/tivoli/products/storage-san-mgr/>

Why Tivoli for Linux

IBM Tivoli Software on Linux helps companies to embrace the future by allowing them maximum flexibility in their choice of new technologies and maximizing company ROI. Tivoli software on Linux provides management capabilities and industry best practices for companies' intelligent systems. This helps to speed time to deployment, maximize business efficiently, and improve the bottom line.

In particular, Tivoli on Linux provides the following benefits:

- ▶ **Reduces the TCO:** Linux helps lower the TCO. Tivoli lowers the cost of operations.
- ▶ **Enhances security:** Linux is attractive but has occasional security concerns. Tivoli helps address these concerns.
- ▶ **Meets service requirements:** Service-level expectations do not decrease with Linux. Tivoli enhances your ability to meet or exceed service-level requirements.
- ▶ **Offers advanced management:** As Linux moves from a Web serving and print serving solution to a critical business applications platform, Tivoli assists in managing the IT resources and the business applications.
- ▶ **Support heterogeneous platforms:** Linux may not be the only platform in your business. Tivoli offers consistent management and licensing across heterogeneous platforms.
- ▶ **Adheres to open standards:** Tivoli software from IBM supports the most open standards. This is an important consideration for Linux-based and other

open source deployments. The end result is unprecedented interoperability between hardware and software components of your infrastructure.

For further information about Tivoli products on Linux, see:

<http://www.ibm.com/software/os/linux/software/tivoli.shtml>

Tivoli Linux support

Tivoli's current strategy is to run under Red Hat, SuSE, and Turbo Linux running kernel 2.4 and server architectures of the xSeries and the zSeries. Tivoli has issued a statement of direction to provide products on the iSeries and pSeries servers in the future.

9.5.1 Tivoli product availability on Linux

The Tivoli brand of products on Linux includes:

- ▶ IBM Tivoli Directory Integrator
- ▶ IBM Tivoli Directory Server
- ▶ IBM Tivoli Access Manager for e-business (formerly known as Policy Director/Tivoli Policy Director for Application Servers)
- ▶ IBM Tivoli Access Manager for Operating Systems (formerly known as Policy Director for Operating Systems (function of Tivoli Security Manager))
- ▶ IBM Tivoli Analyzer for Lotus Domino
- ▶ IBM Tivoli Business Systems Manager
- ▶ IBM Tivoli Configuration Manager
- ▶ IBM Tivoli Decision Support for OS/390
- ▶ IBM Tivoli Distributed Monitoring
- ▶ IBM Tivoli Enterprise Console
- ▶ IBM Tivoli Identity Manager
- ▶ IBM Tivoli Inventory
- ▶ IBM Tivoli Management Framework
- ▶ IBM Tivoli Monitoring for Transaction Performance
- ▶ IBM Tivoli Monitoring (formerly known as Tivoli Distributed Monitoring)
- ▶ IBM Tivoli Monitoring for Databases
- ▶ IBM Tivoli Monitoring for Messaging and Collaboration
- ▶ IBM Tivoli Monitoring for Web Infrastructure (for Apache HTTP Server, IBM WebSphere Application Server, WebLogic Server)
- ▶ IBM Tivoli NetView
- ▶ IBM Tivoli NetView for z/OS (MSM Agents, NMC Server)
- ▶ IBM Tivoli Point of Sale Manager
- ▶ IBM Tivoli Remote Control
- ▶ IBM Tivoli Risk Manager
- ▶ IBM Tivoli SANergy®
- ▶ IBM Tivoli SecureWay® Policy Director

- ▶ IBM Tivoli Service Level Advisor
- ▶ IBM Tivoli Software Distribution
- ▶ IBM Tivoli Storage Manager
- ▶ IBM Tivoli Storage Manager for Enterprise Resource Planning (formerly known as Data Protection for ESS for R/3)
- ▶ IBM Tivoli Storage Manager for Mail (formerly known as Data Protection for Domino)
- ▶ IBM Tivoli Storage Resource Manager
- ▶ IBM Tivoli Switch Analyzer
- ▶ IBM Tivoli System Automation for Linux
- ▶ IBM Tivoli User Admin
- ▶ IBM Tivoli Workload Scheduler

For the most current availability and detailed descriptions of IBM Tivoli Software products for Linux, download the IBM Linux Software Matrix from:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

9.5.2 Tivoli resources

This section provides resources for Tivoli.

Evaluation copy of Tivoli Access Manager

Free evaluation copies of certain Linux-enabled IBM middleware products, complete with installation, configuration, and support documentation are available through the Web.

You can download an evaluation copy of the Tivoli Access Manager, as well as other Linux-enabled IBM Software, from the following Web site:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

Tivoli Developer Domain

IBM provides a valuable set of online technical developer resources through the Tivoli Developer Domain Web site. This Web site, which often includes new technology previews as well as developer experiences, is located at:

<http://www.ibm.com/developerworks/tivoli/>

Relevant Linux-related IBM Redbooks

The current list of relevant Tivoli on Linux Redbooks includes:

- ▶ *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416

- ▶ *IBM Tivoli Monitoring Version 5.1: Advanced Resource Monitoring*, SG24-5519
- ▶ *A Practical Guide to Tivoli SANergy*, SG24-6146
- ▶ *Linux System Administration and Backup Tools for IBM @server xSeries and Netfinity*, SG24-6228
- ▶ *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554
- ▶ *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844
- ▶ *Using IBM Tivoli System Automation for Linux*, REDP-3716
- ▶ *Tivoli NetView Performance Monitor for TCP/IP 1.4 Using Linux as the Web Server*, REDP3602
- ▶ *Securing Linux for zSeries with a Central z/OS LDAP Server (RACF®)*, REDP-0221

The list of Linux-related IBM Redbooks and Redpapers is continually being updated. You can find these documents on the Web at:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

9.6 Lotus products on Linux

Lotus software sets the standard for truly innovative software and services that reflect the brand's unique representation of the new ways that individuals and businesses work together to achieve success. Lotus software is further redefining the concept of conducting business through practical knowledge management, e-business, and other ground breaking ways to connect ideas, thinkers, buyers, sellers, and communities around the world through the Internet. Lotus software is marketed in more than 80 countries or regions worldwide through direct sales and extensive business partner channels.

For detailed information about Lotus products, see:

<http://www.lotus.com>

9.6.1 Lotus product availability on Linux

The following Lotus products are available on Linux:

- ▶ IBM Lotus Web Access (iNotes™)
- ▶ IBM Lotus Domino
- ▶ IBM Lotus Domino for Linux on zSeries
- ▶ IBM Lotus Enterprise Integrator® for Domino

- ▶ IBM Lotus Workflow™
- ▶ IBM Lotus Domino Collaboration Express
- ▶ IBM Lotus Domino Utility Express

For the most current availability and detailed descriptions of IBM Software products for Linux, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

Lotus Domino server

The Domino family provides a multiplatform foundation for collaboration and e-business. It drives solutions from corporate messaging to Web-based transactions and everything in between. Domino is the leading messaging and collaborative application platform available for Linux today. Domino for Linux provides browser-based access for collaborative applications and messaging. Post Office Protocol 3 (POP3), Hypertext Transfer Protocol (HTTP), Internet Message Access Protocol (IMAP), and Lightweight Directory Access Protocol (LDAP) access are all available.

Fully functional Domino solutions are available now on Linux. The key solutions of Domino for Linux involve:

- ▶ Sales force automation
- ▶ Project management
- ▶ Customer relationship management
- ▶ Interactive Web-based forms
- ▶ Resource management

Why Domino for Linux

When you invest in Domino for Linux, you gain:

- ▶ Stability of the operating system
- ▶ Lower cost of ownership (RAS)
- ▶ Powerful remote management
- ▶ An alternative to Microsoft Exchange servers on Windows operating systems
- ▶ High-level of interest in migrating from a Domino on Windows NT environment to Domino on Linux environment

Note: SuSE sells a product called the *SuSE Linux Groupware Server* with Lotus Domino for Intel and AMD systems. It comes prepackaged with Domino R5. You can learn more about this offering on the Web at:

http://www.suse.de/en/company/press/press_releases/archive02/groupware.html

Other Lotus products

Lotus is ready to use the WebSphere portal-based delivery platforms for several Lotus products that are not currently available for Linux. Contact your local Lotus sales representative for additional information.

Lotus has announced plans to deliver the iNotes Web Access product for browsers other than Microsoft's Internet Explorer. These browsers are popular with Linux users.

There are no plans to port the Lotus SmartSuite® desktop productivity products to Linux.

9.6.2 Lotus resources

This section lists various resources for Lotus.

Evaluation copy of Lotus Domino

Free evaluation copies of certain Linux-enabled IBM middleware products — complete with installation, configuration, and support documentation are available through the Web. You can download an evaluation copy of Lotus Domino, and other Linux-enabled IBM Software, from the following Web site:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

Lotus Developer Domain

IBM provides a valuable set of online technical developer resources through the Lotus Developer Domain Web site. This Web site, which often includes new technology previews as well as Lotus developer experiences, is located at:

<http://www.lotus.com/ldd>

Understanding Domino for Linux

This documentation provided with Lotus Developer Domain is a white paper that explores the ways to make the best use of the Linux operating system for Lotus

Domino servers. This white paper also provides important information for Lotus Domino administrators. You can download the white paper from:

[http://doc.notes.net/uafiles.nsf/docs/dlinux/\\$File/dlinuxbasics.pdf](http://doc.notes.net/uafiles.nsf/docs/dlinux/$File/dlinuxbasics.pdf)

Relevant Linux-related IBM Redbooks

Relevant Lotus on Linux Redbooks and Redpapers include:

- ▶ *IBM Lotus Domino 6.5 for Linux on zSeries implementation*, SG24-7021
- ▶ *Lotus Domino R5 for Linux on IBM Netfinity Servers*, SG24-5968
- ▶ *Lotus Domino 6 for Linux*, SG24-6835
- ▶ *Online Briefing: Mastering Domino for Linux*, REDP0039
- ▶ *Domino Web Access 6.5 on Linux*, SG24-7060

You can find a complete list of Linux-related IBM Redbooks and Redpapers on the Web at:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

9.7 Rational Software products on Linux

Rational Software helps organizations create business value by improving their software development capability. The Rational platform encourages teams to:

- ▶ Adopt iterative development practices that reduce project risk
- ▶ Focus on architecture to develop more resilient systems
- ▶ Effectively manage change and protect critical strategic assets

The Rational solution supports multiplatform development environments to ensure interoperability across teams and disciplines. These solutions integrate seamlessly with IBM WebSphere as well as third-party integrated development environments (IDEs) to provide an intuitive design-through-deployment experience.

For detailed information about Rational products, see:

<http://www.ibm.com/software/rational/>

The Eclipse project

This project is an open source software development project dedicated to providing a robust and fully-featured industry platform for the development of highly integrated tools.

Rational has taken a leadership role in the growth of Eclipse as an open source standard that will enable teams to work more effectively on many languages and in many operating environment.

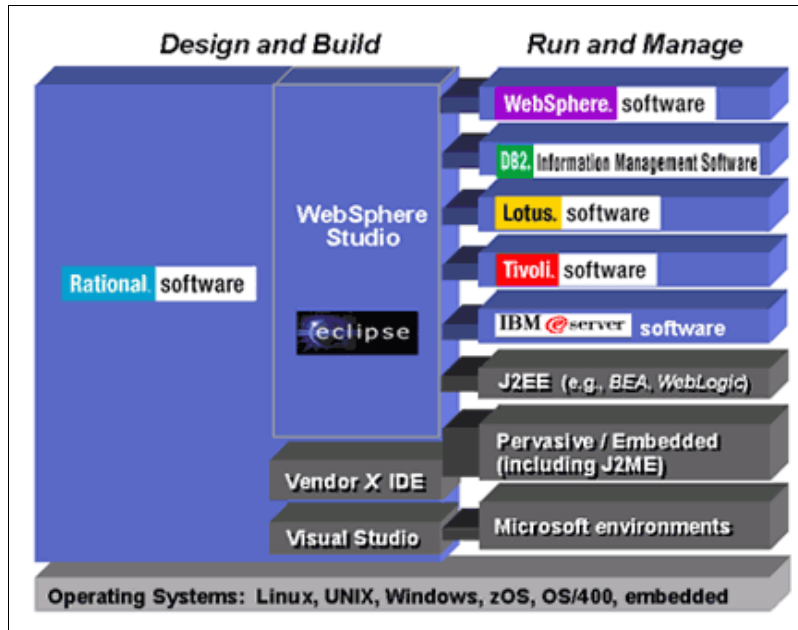


Figure 9-5 Rational platform

For detailed information about the Eclipse project, see:

<http://www.rational.com/solutions/eclipse/>

9.7.1 Rational product availability on Linux

The following IBM Rational products are available on Linux:

- ▶ **Rational ClearCase®:** This product provides software asset management (SAM) for medium to large size teams.
- ▶ **Rational ClearCase MultiSite®:** This product provides automatic, error-free replication for geographically distributed sites.
- ▶ **Rational ClearQuest®:** This product provides activity-based change and defect tracking. It can work with Rational ClearCase to provide a complete SCM solution.
- ▶ **Rational PurifyPlus™:** This product is a runtime analysis solution designed to help developers write faster, more reliable code.

- ▶ **Rational Rose®:** This product is a model-driven development with the Unified Modelling Language (UML), targeted for commercial products and systems.
- ▶ **Rational Test RealTime:** This product provides solution for component testing and runtime analysis.

For the most current availability and detailed descriptions of IBM Rational Software products for Linux, download the IBM Linux Software Matrix from:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

9.7.2 Rational resources

This section lists a variety of resources for Rational Software.

Software development downloads

Free trial of certain Linux products for software development, complete with support, is available on the IBM Internet site. You can download an evaluation copy of Rational products, and other software development tools from the following Web site:

http://www14.software.ibm.com/webapp/download/category.jsp?s=c&S_TACT=&S_CMP=&cat=swdev

Rational Developer Domain

IBM provides a valuable set of online technical developer resources through the Rational Developer Domain Web site. This Web site, which often includes new technology previews and Rational developer experiences, is located at:

<http://www.ibm.com/developerworks/rational>

Introducing the PurifyPlus family

This white paper is a starting point to help understand the tools in the PurifyPlus family. This paper provides an introduction to the practice of runtime analysis and an in-depth comparison of the tools in the family. You can see and download the white paper from the following Web site:

<http://www.ibm.com/developerworks/rational/library/665.html>

9.8 Additional IBM Software

In addition to primary brand software, IBM has made several of its other software products available on Linux.

9.8.1 Product availability on Linux

Non-specific brand IBM Software for Linux includes:

- ▶ AFS® (Andrew File System)
- ▶ Application Workload Modeler for Linux
- ▶ IBM Bank Teller Business Components
- ▶ IBM BookManager® Bookserver
- ▶ IBM Communications Server for Linux
- ▶ IBM Cluster Systems Management (CSM) for Linux
- ▶ IBM Developer Kit for Linux
- ▶ IBM Director
- ▶ IBM Everyplace Wireless Gatekeeper
- ▶ IBM General Parallel File System
- ▶ IBM HTTP Server
- ▶ iSeries ODBC Driver for Linux
- ▶ IBM LANDP®
- ▶ IBM Object REXX
- ▶ IBM Screen Customizer
- ▶ StreamServer Business Communications Platform
- ▶ IBM VisualAge Generator Server
- ▶ IBM VisualAge Pacbase Linux Server
- ▶ IBM VisualAge Smalltalk Enterprise
- ▶ IBM Workstation APL2® for Multiplatforms

For current availability and detailed descriptions of IBM Software products for Linux, download the IBM Linux Software Matrix from:

<ftp://ftp.software.ibm.com/software/linux/IBMSoftwareOnLinux.pdf>

9.8.2 Additional resources

Consult the following resources for more information.

Evaluation copy of IBM Directory Server

Free evaluation copies of certain Linux-enabled IBM middleware products, complete with installation, configuration, and support documentation are available through the Web. You can download an evaluation copy of IBM Directory Server and other Linux-enabled IBM Software from the Web at:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

IBM Redbooks

For more information, consult the following IBM Redbooks:

- ▶ *Linux Clustering with CSM and GPFS*, SG24-6601
- ▶ *Building a Linux HPC Cluster with xCAT*, SG24-6623

9.9 IBM Integrated Platform for e-business

The IBM Integrated Platform for e-business provides intensive testing, one-stop shopping, and a single point of contact for support for the IBM core e-business products and SLES. This combination allows you to:

- ▶ Deploy advanced e-business applications, capitalizing on the latest industry standards such as Java, WebSphere, Linux, and IBM @server technology.
- ▶ Integrate flexible solutions with traditional workloads and mission-critical databases.
- ▶ Save time and resources to design the infrastructure. It is already created, configured, and tested.
- ▶ Tailor the system to customer needs using the breadth of IBM knowledge and experience in e-business.
- ▶ Position customer business for the future IBM @server with the latest advances while leveraging existing technology.
- ▶ Keep TCO low by managing a customer's business centrally and exploiting the economies of server consolidation.
- ▶ Benefit from a total solution for e-business integration.

The Integrated Platform for e-business is currently available for all IBM @server platforms. Table 9-2 shows the software and support included in the offering.

Table 9-2 Components available per platforms

Components	xSeries	zSeries	iSeries	pSeries
Web server	IBM HTTP Server	IBM HTTP Server	IBM HTTP Server	IBM HTTP Server
Application server	WebSphere Application Server	WebSphere Application Server	WebSphere Application Server	WebSphere Application Server

Components	xSeries	zSeries	iSeries	pSeries
Database	DB2 UDB (option)	DB2 UDB (option) or DB2 Connect (optional for component for access to DB2 for z/OS)	DB2 for OS/400	DB2 UDB (option)
Directory (LDAP)	IBM Directory Server	IBM Directory Server	IBM Directory Server	IBM Directory Server
Operating system (delivered by SuSE)	SLES	SLES	SLES	SLES
Server hardware	Models 305, 335, 345, 360, and 440	Models 800, 900, and 990	Models 810 and 825	Models 630 and 650
Storage	Integrated HDDs with RAID-1 and optional RAID-5 EXP300 also an available option	ESS storage option available	Integrated iSeries storage	Integrated HDDs with RAID-1 and optional external storage with RAID-5
Firewall (required for Internet topology only)	Optional third-party component	Optional third-party component	Netfilter/iptables Optional third-party component	Optional third-party component

IBM Integrated Platform Express is also available for the xSeries platform. You can find detailed information about the Integrated Platform for e-business on the xSeries server on the Web at:

<http://www.ibm.com/servers/solutions/linux/integrated/xseries/index.html>

For detailed information about the Integrated Platform for e-business on zSeries, see:

http://www.ibm.com/servers/eserver/zseries/os/linux/integrated/index.html?ca=zseries&me=W&met=drive_from_eserver_IP

For detailed information about the Integrated Platform for e-business on iSeries, see:

<http://www.ibm.com/servers/eserver/iseries/linux/ipeb/>

For detailed information about the Integrated Platform for e-business on pSeries, see:

<http://www.ibm.com/servers/eserver/pseries/linux/integrated/>

9.10 IBM Software on Linux deployment examples

The IBM Solution Scenario Profile Database (SSPD) is a unique reuse mechanism that contains synopses of both solutions tested by IBM development labs and solutions created for clients. Each entry describes the problem being solved and how the pieces fit together. It also includes a diagram, list of products used, and supporting material.

Use the fully searchable database to find which levels of products were tested together. You can also find details on specific solutions from which you can base your solution.

The following sections offer a brief look into examples taken from the SSPD. For detailed information about the Solution Scenario Profile Database, see:

<http://sspd.ibm.com/>

Restriction: This is an IBM intranet site that is available only for IBM representatives.

9.10.1 Using WebSphere Application Server on the zSeries server

A client needed a platform-independent Revenue Accounting System (RAS) application to provide general access from the Internet and employee access from the intranet and Internet. They also needed to modernize their internal-only application called the *Flight Schedule Inquiry System* to provide employees both intranet and Internet access to improve their ability to check work schedules from anywhere in the world.

The solution was in using IBM WebSphere Application Server, Advanced Edition V3.5; WebSphere Studio, Advanced Edition; VisualAge for Java; and Linux for zSeries. This allowed the client to quickly update business applications, consolidate workloads onto IBM mainframes from a variety of servers, realize savings, and increase availability and reliability.

This solution employed the following IBM products:

- Hardware:
 - zSeries 900

- zSeries running Linux
- ▶ Software:
 - WebSphere Application Server, Advanced Edition
 - VisualAge for Java

For a full description and diagrams of this scenario, see:

<http://w3-1.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/TD100650>

9.10.2 Using Domino and WebSphere on xSeries and iSeries

A client who offered hosting services was experiencing problems with server performance. Since they needed to handle many operating systems to support their customers' needs, they found it difficult with their Windows 2000 Domino systems.

To achieve their goals for upgrading their environment, this client invested in several xSeries and iSeries servers and Linux to support their internal and hosting applications. Linux allowed this client to work with their customers who wanted to test and run production middleware, as well as maximize cost efficiency. Overall, this solution allowed the client to reduce costs, increase performance, and offer a better and more affordable service.

This solution employed the following IBM products:

- ▶ Hardware
 - xSeries 200 and 220
 - iSeries 270
- ▶ Software
 - Lotus Domino
 - Apache
 - Lotus Sametime®
 - QuickPlace®
 - DB2 Version 7.1
 - WebSphere Application Server, Advanced Edition
 - Lotus LearningSpace® Core
 - Firewall
 - Logical partition (LPAR)

For a full description and diagrams of this scenario, see:

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/TD100883>

9.10.3 DB2 UDB, DB2 Connect, and WebSphere Portal

A client developed and deployed its first Web implementation on non-IBM Windows NT hardware. This included two applications for member stores to reorder items through the Web. The client planned to develop many more applications, but was concerned with the ability of its Windows NT environment to support them.

The client turned to IBM for help. They converted their existing Windows NT applications to zSeries servers running SuSE Linux V7. They consolidated their two applications from 40 dual-processor Windows NT machines for development and production to run on two zSeries LPARs. The software environment now uses WebSphere Application Server, Advanced Edition, Version 3.5 on Linux to run custom applications for their e-commerce needs. Developers use WebSphere Commerce Studio V5.1, which includes VisualAge for Java V3.5, to develop the new applications and deploy them to Linux for zSeries. And, the client's users migrated from Microsoft and Lotus cc:Mail to Lotus Notes.

The client used WebSphere Commerce Suite and WebSphere Application Server (both on Linux for zSeries) to implement an initiative for a robust portal for independent grocers, distributors, and retail stores. This plan allows these groups to access important business-critical information through the portal. This includes ordering products, finding current pricing, and conducting forms management, all within the portal environment.

This solution employed the following IBM products:

- ▶ Hardware
 - zSeries
- ▶ Software
 - DB2, DB2 UDB, and DB2 Connect
 - IBM HTTP Server
 - Lotus Domino, Lotus iNotes, Lotus QuickPlace, and Lotus Sametime
 - VisualAge for Java
 - WebSphere Application Server
 - WebSphere Application Server, Advanced Edition
 - WebSphere Application Server, Advanced Edition Version 3.5
 - WebSphere Commerce Studio
 - WebSphere Commerce Suite
 - WebSphere Commerce Studio Version 5.1

For a full description and diagrams of this scenario, see:

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/TD100764>

9.10.4 Business continuity solution for Disaster Recovery

This client's enterprise environment had a failover site for the zSeries servers, but none for the pSeries server and RS/600 H70. When a typhoon hit in 2001, the company experienced great damage. It was unable to operate the core applications running on its pSeries machine due to the power outage, and it became clear that having a failover site for the pSeries environment was as important as having a failover site for the zSeries environment.

Meanwhile, keeping the best service and performance while adhering to a limited budget was another main business priority for the company. The current stock inventory system on the pSeries server couldn't handle the growing business requirements unless the server memory and disk space could be increased by 20%. It was clear something had to change. Because the company was satisfied with the functionality of its IBM @server environment, it looked for help from IBM. The client realized that IBM was the only vendor who understood its environment and had the capability and expertise to meet its requirements. Thus, they charged IBM with balancing the transaction workload and reducing the cost of hardware, software, and application maintenance.

IBM Global Service - Integrated Technology Services, Technical Support Service for zSeries (TSSZ) helped the client to create a solution by leveraging their existing servers and hardware. This solution employed the following IBM products:

- ▶ Hardware
 - zSeries mainframes (9672-R26 and R15)
 - A tape library
 - Enterprise Storage Server (ESS)
 - pSeries server
- ▶ Software
 - IBM Tivoli Storage Manager client based on Turbo Linux for S/390 V6.5
 - IBM Tivoli Storage Manager server based on OS/390 V2R10
 - removable media (RMM) based on OS/390 V2R10
 - Extended Remote Copy (XRC) based on OS/390 V2R10
 - peer-to-peer remote copy (PPRC) based on OS/390 V2R10.

The implementation was completed in three steps:

1. IBM and ACU Corporation, an IBM Business Partner, worked together to migrate the stock inventory system from pSeries to zSeries. The COBOL applications previously ran on the pSeries servers and were ported to Linux running on the zSeries. This leveraged the existing COBOL compiler and runtime environment to reduce the labor costs and the time required for language translations.

2. After total migration of the pSeries applications to Linux, the team implemented IBM Tivoli Storage Manager. The IBM Tivoli Storage Manager client provided a good storage solution by combining IBM Tivoli Storage Manager with a new removable media (RMM) tape library. The backup server was consolidated onto the zSeries platform, which made it easier to backup and restore data, regardless of the data's format.
3. XRC/PPRC was implemented to help the client create a failover site to deal with disaster and recovery. By using the resources on another existing zSeries mainframe, the failover site can take over the main site's job without much effort. The potential impact to the client would be reduced bandwidth from having to transfer data through XRC/PPRC feature, which would cause network traffic to increase. After evaluating and discussing the plan, the client decided the load was acceptable. Since the client has been a user of a zSeries mainframe platform for years, it was familiar with this mainframe architecture. The solution currently has about 200 users.

After the applications were ported to Linux on zSeries, the client could redeploy the original pSeries servers, thereby protecting its pSeries investment. And, by moving to zSeries, the client benefits from better computing power, substantial performance improvement, and a resulting cost reduction of 30%. In addition, the new tape library can be fully used, which results in up to a 30% reduction in storage costs and frees current direct access storage device (DASD) space.

For a full description and diagrams of this scenario, see:

<http://w3-1.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD100867>

9.11 IBM and Linux solution customer success

Several clients have experienced great success in implementing an IBM and Linux solution for their diverse needs. Here are a few of those clients and links to where you can learn about their stories:

► **WebSphere: Diveo Broadband Networks, Inc.**

Diveo Broadband Networks transitioned to a WebSphere for Linux architected solution. In doing so, they achieved increased affordability and reliability, increased availability, and savings on maintenance costs.

For more information about this success story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-4Z8VXP>

► **DB2: Grede Foundries, Inc.**

Grede Foundries transitioned to a DB2 for Linux architected solution. This brought them improved reliability, scalability, performance, stability, and cost savings.

For details about this success story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-4XHRGT>

► **Tivoli: Tamkang University**

Tamkang University deployed a Linux-based solution. This solution automated backup and recovery, and optimized storage allocation. It also gave them improved reliability, capacity, and performance.

For details about this success story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5942EY?OpenDocument&Site=software>

9.12 Additional IBM Software for Linux resources

IBM provides its employees, partners, and clients with an array of resources designed to assist and improve the deployment of Linux solutions. This section highlights some of those resources.

IBM Software on Linux: Customer Reference (GK3T-4174-02)

The customer resource CD has success stories, testimonials, reports, articles, brochures, and other information about IBM Software on Linux.

Version date: 01-09-2004

SMB Linux sales kit for BPs (GK3T-4176-00)

The customer resource CD has success stories, testimonials, reports, articles, brochures, and other information about SMB for Business Partners.

Version date: 10-07-2003

IBM Software on Linux brochures

- IBM-Software-Linux, GC27-1498
- IBM WebSphere for Linux, GC27-1692
- IBM DB2 for Linux, GC27-1690
- IBM Lotus for Linux, G325-2420
- IBM Tivoli for Linux, GC27-1689

9.12.1 Evaluation and demonstration software

To get started with IBM Software Group products, the following sections provide information about the tools that are available for evaluation or download.

IBM Speed Start Your Linux Application program

In an effort to provide hands-on experience with IBM middleware on Linux, IBM provides customer evaluation copies of several of its major software products for Intel-based Linux systems.

Among the products currently available are:

- ▶ IBM DB2 UDB Enterprise Server Edition (ESE) V8.1 for Linux
- ▶ IBM DB2 Migration Toolkit V1.1 for Linux
- ▶ IBM Directory Server V5.1 for Linux
- ▶ IBM Lotus Domino 6 for Linux (English)
- ▶ IBM Tivoli Access Manager Base 4.1 for Linux
- ▶ IBM WebSphere Application Server V5 for Linux
- ▶ IBM WebSphere Studio Site Developer V5 for Linux
- ▶ IBM WebSphere MQ V5.3.1 for Linux
- ▶ IBM WebSphere SDK for Web Services for Linux

You can find a description of these products and links to freely available evaluation software and installation documentation on the Web at:

<http://www.ibm.com/developerworks/offers/linux-speed-start/>

Free technical support is available for these applications through a Web forum at:

http://www.ibm.com/developerworks/forums/dw_forum.jsp?forum=160&cat=5

The Linux Software Evaluation Kit

If you are downloading evaluation software from the Speed Start Your Linux Application Web site is too cumbersome, you can obtain the software on DVD. It comes as a set of two DVDs entitled *Linux Software Evaluation Kit (SEK)*. The Linux SEK is limited to Intel-based Linux implementations. It is free of charge and includes limited Web forum technical support.

For detailed about the Linux Software Evaluation Kit, see:

<http://www.ibm.com/developerworks/offers/linux-speed-start/>

Additional Linux-enabled software for free download

IBM provides a large number of trial, beta, and free software products for Linux that you can download from the Web. You can find a complete list of this IBM Linux-enabled software, which is available for download free of charge at:

<http://www14.software.ibm.com/webapp/download/home.jsp?s=p>

9.12.2 IBM Software on Linux training

IBM Learning Services offers an extensive catalog of world class training courses. Its portfolio of training services includes on-site training, consulting, development, delivery, transformation, learning management, and outsourcing.

The courses about IBM Software on Linux include:

- ▶ **DB2 for Linux:** CF201 - *DB2 Admin Workshop for Linux*
http://www.ibm.com/services/learning/MainServlet.wss?pageType=course_description&courseCode=CF201
- ▶ **WebSphere on Linux:** SW247 - *WebSphere System Administration on Linux*
<http://www.software.ibm.com/wsdd/education/enablement/curriculum/sw247.html>

Additional courses are being prepared. For an updated list of Linux courses offered by IBM Learning Services in the United States, see:

<http://www.ibm.com/services/learning/us/catalog/linux/>

9.12.3 IBM Redbooks

IBM Redbooks are written and published by the IBM International Technical Support Organization (ITSO). The ITSO develops and delivers skills, technical know-how, and materials to technical professionals including IBM, Business Partners, clients, and the marketplace. The ITSO value-add information products address product, platform, and solution perspectives. They explore integration, implementation, and operation of realistic customer scenarios.

The ITSO partners with IBM Divisions, Lotus Development, and Tivoli Systems in the process of IBM Redbook development. The ITSO is part of the IBM Global Technical Support organization within IBM Global Sales and Distribution.

IBM Redbooks, RedDrafts, and Redpapers are available to the public for printed copy purchase or freely downloadable as PDF files. For Linux-related Redbooks, see:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

9.12.4 IBM Software support organizations

IBM has established a core set of organizations to proof and improve client solutions based on IBM Software, as well as contribute to the evolution of IBM middleware products.

IBM Linux Integration Centers

The Linux Integration Centers (LICs) support and propagate the use of IBM middleware products for Linux solutions. These presales support centers are located in Austin (Texas), Boeblingen (Germany), Singapore, Beijing (People's Republic of China), and Tokyo (Japan). They provide assistance to a variety of industry clients including financial services, retail, and manufacturing in order to identify key areas within a client's infrastructure to implement Linux. The LICs provide IBM hardware, software, and services. They can also simulate computing environments to provide detailed consultations, architectural guidance, and educational briefings surrounding Linux and open source software.

Skilled in a broad spectrum of operating systems, LIC team members can quickly prepare an assessment of a client's infrastructure and develop a step-by-step blueprint for transitioning to a Linux environment on IBM Software platforms. The team members include system architects, database administrators, project managers, and operating system specialists.

The LIC teams provide the following valuable services:

- ▶ Assist installation, configuration, and product integration during Proof of Concepts
- ▶ Guide architectural designs by incorporating best practices as well as using the experience gained from prior Linux migrations
- ▶ Find solutions to client problems by replicating the client's environment and debugging problems
- ▶ Interface with the owner of each product to assess requirements and resolve issues
- ▶ Provide engagement-based support for qualified Linux Middleware opportunities by offering:
 - Technical consulting

- Product integration
- Deployment planning
- ▶ Drive Linux as an e-business application server
- ▶ Integrate middleware and applications (ISV or customer)
- ▶ Document migration paths
- ▶ Identify supported hardware and software stack environments
- ▶ Gather performance benchmarks
- ▶ Demonstrate data conversion
- ▶ Create and leverage reusable intellectual assets

To qualify as a candidate for assistance from a Linux Integration Center, a sales opportunity must:

- ▶ Have significant SWG middleware content
- ▶ Allow the LIC to build collateral or assets for reuse
- ▶ Have the potential for a customer to become a referenceable account
- ▶ Support current IBM or SWG strategic initiatives
- ▶ Be an xSeries or zSeries opportunity with significant SWG middleware
- ▶ Include migration from SUN to Linux and IBM middleware
- ▶ Include migration from Oracle to DB2
- ▶ Include migration from BEA WebLogic to IBM WebSphere
- ▶ Have complete OMNotes Entry identified with LINX opportunity code

Worldwide Linux Advocates

IBM has developed the Linux Advocates program as a complement to the Linux Integration Center. If a Linux sales opportunity does not meet the requirements to be a candidate for Linux Integration Center (LIC) assistance, consider a Linux advocate for engagement.

A Linux Advocate is an IT Architect or IT Specialist who is level 4 skilled in IBM Software brand products (DB2, WebSphere, Lotus, or Tivoli) on Linux.

The Linux Advocate is a valuable resource by:

- ▶ Assisting the technical and sales team with Linux-based IBM Software product technical questions and issues
- ▶ Influencing the technical and sales community so that more Linux-based solutions are recommended to our clients
- ▶ Providing Linux leadership on their assigned accounts
- ▶ Being up-to-date on the evolution of Linux, especially as it relates to the deployment of IBM Software for customer solutions

Linux advocates are assigned in every worldwide geographic region. For a updated list of Linux advocates, refer to your sales contact.

Global e-business Solutions Center

The Global e-business Solution Center (GeSC) is a part of the IBM Global Technical Sales Support team. Its mission is to provide architectural and ISV solutions integration technical sales support leadership to help IBM teams win and clients to succeed with integrated e-business solutions.

The GeSC provides complex technical sales support for:

- ▶ End-to-end solution architecture and ISV integration solution selling
- ▶ Scalability and performance of cross-platform end-to-end solutions

To carry out this mission, the GeSC focuses on the following major support approaches:

- ▶ Complex client sales engagements
- ▶ Sales call technical assistance
- ▶ Solution and architecture workshop
- ▶ Solution proof of concept
- ▶ Solution scalability and performance assessment

As part of this responsibility, the GeSC provides proactive second-level support directly to the Sector/Solutions Sales Teams in the Americas. For EMEA and Asia Pacific, the GeSC is focused on leveraging its capabilities to help build geography-based, strong technical sales support skills, and back them up on critical, high-impact customer situations.

Solution Partnership Centers: Third-party application porting

Ten worldwide Solution Partnership Centers (SPCs) offer easy on-site and remote access for developers to IBM resources. SPCs are state-of-the art facilities equipped with the latest IBM @server and software technologies, staffed by IBM technical consultants.

To learn more about IBM Solution Partnership Centers, see:

<http://www.developer.ibm.com/spc/index.html>

Linux Competence Center

The Linux Competence Center (LCC) in Poughkeepsie, New York, is the result of close cooperation between the IBM eServer Briefing Center, the New Technology Center, and the eTP Design Center. While the LCC is focused primarily on Linux under zVM server consolidation solutions, it can provide

demonstrations of WebSphere with servlets connecting to CICS, DB2, WebSphere MQ, and IMS.



Linux clusters

A cluster is a collection of interconnected complete systems that work together to provide a single, unified computing capability. The resulting cluster can either be a parallel system or a distributed system.

This chapter discusses:

- ▶ Clusters
- ▶ Linux cluster components
- ▶ Linux cluster history
- ▶ Linux cluster solutions
- ▶ The future of Linux clusters

10.1 An introduction to clusters

This section reviews the major cluster concepts and types. It covers the benefits of a cluster, and the role that clusters play in emerging computer trends.

10.1.1 Core cluster concepts

As mentioned earlier, a cluster can be roughly defined as a system (either parallel or distributed) that is built from interconnected servers. It is viewed as a single entity to application or user resources.

The following sections explain the benefits, and the traditional and emerging roles of clusters.

Cluster benefits

Clusters are justified by several advantages:

- **Redundancy:** Multiple physical cluster nodes avoid the fatal flaw of traditional massively parallel processing (MPP) and symmetric multiprocessing (SMP) systems. In these systems, multiple processors are stored in a single physical computer creating a single point of failure. Additionally, because cluster nodes are traditionally based on commodity hardware, the replacement cost for a failed node is typically less than repairing a similar component on an MPP system.

Figure 10-1 shows an example of an SMP solution where processors and other computing resources are grouped in a single package.




Figure 10-1 SMP system with shared storage and memory

Figure 10-2 shows an example of a grouping of separate servers in an MPP arrangement.




Figure 10-2 MPP system with an individual processor and memory

- **Scalability:** Clusters can scale into the low thousands of heterogeneous nodes while remaining competitive on price and performance against MPP systems.
- **Availability:** Because clusters have multiple interconnected physical nodes and resources (as shown in Figure 10-3), they can ensure better data continuity and recovery operations.




Figure 10-3 An availability cluster

- **Resource sharing:** Clusters can be the optimum solution for computing environments that need to dynamically assign and manage processing, and storage resources in a multiuser environment.

Traditional cluster roles

Originally, clusters were part of a larger parallel computing heritage that emphasized the speed of calculations above all else. Over time, the use of clusters has expanded into other areas. The following three areas are the traditional computing roles played by clusters.

High Performance Computing (HPC)

HPC clusters are designed to use parallel computing to apply more processor power to the solution of a problem. There are many examples from scientific computing where using multiple low-cost processors in parallel to perform large numbers of operations can be highly cost effective while providing massive amounts of computing power. This is referred to as *parallel computing* or *parallelism*. Table 10-1 shows the application areas and some examples of HPC.

Table 10-1 HPC application areas and examples

Application area	Application examples
CAE (Computational Fluid Dynamics)	<ul style="list-style-type: none"> ► Fluent ► STAR-CD ► PowerFlow (from Enterprise X-Architecture)
CAE (Crash)	<ul style="list-style-type: none"> ► LS-DYNA ► PAM-CRASH ► Radioss
CAE (Structures)	<ul style="list-style-type: none"> ► Nastran ► Ansys ► Abaqus
Petroleum (Reservoir Modeling)	<ul style="list-style-type: none"> ► Eclipse ► VIP
Life Sciences (Ab Initio Chemistry)	<ul style="list-style-type: none"> ► Gaussian ► Gamess ► Jaguar
Life Sciences (Molecular Modeling)	<ul style="list-style-type: none"> ► Charmm ► Amber

Typically, parallelism in a cluster is achieved through message passing (Figure 10-4). Message parsing involves an application running on one node that takes a large set of data, breaks it up into smaller chunks, and passes those chunks to the other nodes in the cluster. As each node finishes processing its

respective chunk, it returns that chunk to the node of origin, which reassembles the information into a result set.

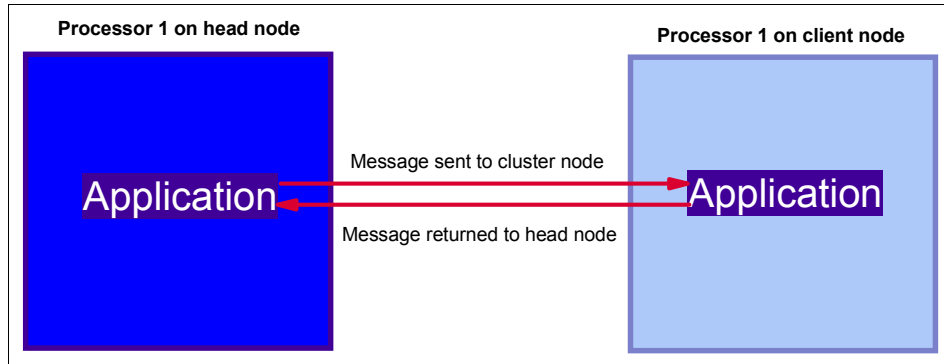


Figure 10-4 A simplified view of cluster message passing

The bottleneck in most clusters is not the processor, but rather the communication infrastructure used to pass messages from processor to processor to include the data bus of each node. The data bus in most commodity computer hardware operates at a noticeably slower speed than the processor (or even the network infrastructure) does. Because of this, clusters typically do not scale much beyond a few thousand nodes, even though the most powerful high performance commodity clusters are usually between 1,024 and 2,048 nodes.

Creating an architecture for this kind of cluster brings its own challenges. These include:

- ▶ How to install and maintain the operating and application environment on all nodes
- ▶ How to proactively manage these nodes through issuing commands and gracefully handling failures
- ▶ The requirement for parallel, concurrent, and high performance access to the same file systems
- ▶ Inter-process communication between the nodes to coordinate the work that must be done in parallel

The goal is to provide the image of a single system by managing, operating, and coordinating a large number of discrete computers. Often in this environment a user interacts with a specific node to initiate or schedule a job to be run. The application, in conjunction with various functions within the cluster then determines how this job is spread across the various nodes of the cluster to take advantage of the resources available to produce the desired result.

High Throughput Computing (HTC)

While science needs have driven HPC and parallel operations (and consequently, cluster technology), some areas of research place a higher emphasis on throughput (quantity) of data operations than on the speed of those data operations.

Note: A common analogy used to explain the difference between HPC and HTC is the fighter jet versus the commercial airliner jet. While a fighter jet can carry a single individual to a destination faster than an airliner jet can, the airliner jet can carry multiple (more than 200) individuals at a time on the average, providing a more efficient transportation medium.

Typically, these types of operations can take weeks or even months of computational time to solve. These operations are also distinguished from HPC from the lower coupling requirements that are needed. A simple definition of coupling is the required level of data sharing among nodes at any one time. Because many HTC environments do not meet the strict definition of parallel computing (either through design or in execution), they are often referred to as *distributed computing environments*.

Common examples of HTC applications include protein folding and radio astronomy, among others.

To learn about protein folding, see:

<http://folding.stanford.edu/>

To learn about radio astronomy, see:

<http://setiathome.ssl.berkeley.edu/>

A variety of HTC applications run on common client operating systems. You can learn about these applications on the Web at:

<http://www.techimo.com/teams.html>

Blue Gene research project: This goal of this IBM research project is to build a massively parallel computer to be applied to the study of biomolecular phenomena such as protein folding. The project was expanded by the addition of Blue Gene/L. Blue Gene/L is also part of IBM's research initiative in autonomic computing to design computer systems that are self-healing, self-managing, and self-configuring. To learn more about this research, see:

http://www.research.ibm.com/resources/news/20031114_bluegene.shtml

Figure 10-5 shows an example of protein folding for a typical HTC application. Some considerations of HTC clusters include:

- ▶ Defining the minimum level of coupling required by the HTC applications
- ▶ Determining how best to allocate computing resources among multiple users of the cluster
- ▶ Resolving conflicts that may arise if resource ownership is distributed (that is, different entities or individuals own their own nodes that participate in the cluster)




Figure 10-5 Protein folding simulation: Typical HTC application

High-availability cluster

High-availability clusters (Figure 10-6) are typically built with the intention of providing a resilient environment through redundancy. They provide a computing environment where the failure of one or more components (hardware, software, or networking) does not significantly affect the availability of the application or applications being used.

In the simplest case, two computers may be configured identically with access to shared storage. During normal operation, the application environment executes on one system, while the other system simply stands by ready to take over running the application in the case of a failure.




Figure 10-6 High availability example: A Web server farm

When a failure occurs, the second system takes over the appropriate resources (storage, networking address, and so forth). This process is typically called *failover*. The second system then completely replaces the failed system. The end users do not need to know that their applications are running on a different physical machine.

As you may expect, actual high availability clusters can be complex to design, especially if shared; coherent data storage is required. In fact, two subcategories of high availability clusters are file storage and database cluster, each with their own particular requirements and demands. Understanding all possible failure types and building in redundancy and the processes required to provide a smooth, efficient, and transparent failover is a difficult endeavor.

Some of the abilities of high availability clusters include:

- ▶ Providing an administrative switch over between nodes for maintenance, upgrade, or other non-failure situations
- ▶ Maintaining a consistent environment between all nodes
- ▶ Recognizing when a failure has occurred
- ▶ Implementing the failover process so that all required resources are available to the standby node
- ▶ Providing transaction and data integrity in case of a failure

All three areas, despite their differences, attempt to provide some level of abstraction that allows users, and to an extent applications, to see the cluster as a single resource instead of individual computers.

Emerging roles

Over time, clusters are being used for more targeted use. Variations on the three traditional cluster roles are discussed in the following sections.

Server consolidation

Clusters can be an ideal platform for server consolidation because they provide a single point of control for software installation, hardware control, system monitoring, and application maintenance. Parallel or Web-based applications can take advantage of the scaling (see “Load balancing” in the following section) and price-performance capabilities of a cluster. Clusters can also make an ideal front end in a multitiered architecture.

Some considerations include:

- ▶ Centralizing management of resources that previously were scattered across multiple sites
- ▶ Greater economy of scale from reducing physical redundancy
- ▶ Data integration provides higher availability, improved recoverability, and easier management of enterprise data
- ▶ Application integration provides an opportunity for business process synergy

Load balancing

Load balancing (and, by extension, horizontal scaling) involves multiple servers providing a single interface to a set of resources that can arbitrarily grow (or shrink) in size over time. The most common example of this is a *Web server farm*. Here, a single Universal Resource Locator (URL) (the interface) is used to access a Web site (the resource). Requests coming in through that interface can be allocated across a large set of servers. This provides higher capacity and the

ability to manage the end user experience through such functions as load balancing.

Figure 10-7 shows an example of how load balancing was used in the 1998 Winter Olympics.

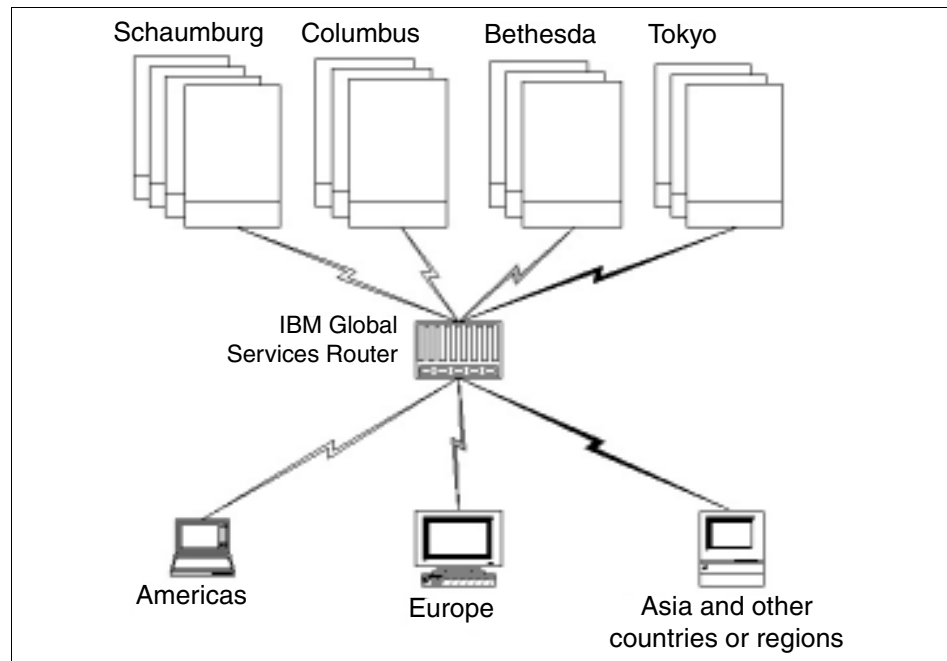


Figure 10-7 Load balancing for the 1998 Winter Olympics

Scaling considerations include:

- ▶ Establishing baseline estimates of high and low resource usage limits
- ▶ Determining which elements of an application must be reworked (if any) to take advantage of a horizontally-scalable cluster

On-demand computing

The technologies that will eventually support on demand computing are still in their infancy. These technologies include grid and autonomic computing. There is little doubt that clusters will have a role to play. See 10.5, “The future of Linux clusters” on page 346, for more information.

10.1.2 IBM cluster offerings

IBM has several product-based cluster offerings. This section discusses some of them.

Parallel sysplex

Each node in the parallel sysplex can have a single CPU or consist of a symmetric multiprocessor. Each node runs a separate copy of the operating system.

The processors in an SMP node run a single operating system, zOS (previously known as MVS) on the IBM S/390. They have shared main memory. The nodes in the parallel sysplex have shared disks, which you can access directly from each node. You can configure large complexes of shared disks, using the Fibre Connection (FICON) switch.

SP

The SP system is a general-purpose scalable parallel system. It is based on *share-nothing architecture*. Share-nothing refers to the fact that each processor has its own cache and disk subsystem.

Generally available SP systems range from 2 to 128 processor nodes. Each processor node is functionally equivalent to a stand-alone RS/6000 workstation or server.

Large SP systems with up to 512 nodes have been delivered and are successfully being used today. Each processor node contains its own copy of the standard AIX operating system, and other standard RS/6000 system software.

A set of new software products designed specifically for SP allows the parallel capabilities of SP to be effectively exploited. The SP system is a state-of-the-art horizontally scalable environment.

Cluster 1600

The IBM Cluster 1600 extends and enhances IBM's innovative, proven AIX and RS/6000 SP clustering technologies while adding support for Linux. The Cluster 1600 includes both the legacy SP system and new clusters made up of pSeries servers.

While the Cluster 1600 still includes the Parallel System Support Programs (PSSP) for AIX 5L software, IBM has introduced the Cluster Systems Management tool as part of an overall “building block” cluster methodology. In this new approach, IBM is building cluster capability into all pSeries hardware, and providing CSM on both AIX 5L and Linux to managed these capabilities.

Note: CSM is also the utility of choice for Cluster 1350, providing a consistent administrative utility across both cluster offerings.

CSM extends the functions of the original PSSP to support the installation, operation, management, and administration of the RS/6000 SP, attached servers, and Clustered Enterprise Servers (CES) from a single point of control.

Note: For more detailed information about IBM Cluster 1600, see:

http://www-1.ibm.com/servers/eserver/clusters/hardware/1600_facts.pdf

10.2 Linux cluster components

Much focus is given to computers (or nodes) that make up a cluster. Just as important are the communication and storage technologies that work to support the cluster, and the management software that keeps the cluster running properly.

The cluster applications themselves provide the high availability and resiliency necessary in a cluster. Throughout this section, we match general cluster component areas with specific components of Cluster 1350.

10.2.1 Nodes

Clusters (especially high performance clusters) typically have a large number of computers. In general, most of the nodes are configured identically. The goal of identical configuration is that the individual tasks that make up a parallel application will run equally well on the node to which they are assigned. However, some nodes in a cluster often have physical and logical differences.

The following sections discuss logical node functions and the physical node types.

Node functions

A *cluster* is two or more (often many more) computers working as a single logical system to provide services. Although, from the outside, the cluster may look like a single system, the internal workings to make this happen can be quite complex. The logical functions that are explained in the following sections are found in most clusters.

These are logical functions. In some cases, multiple logical functions may reside on the same physical node. In other cases, a logical function may be spread

across multiple physical nodes. Do *not* underestimate the *importance* of cluster management. Cluster management is imperative when coordinating the activities of large numbers of systems.

Compute node

The compute node is where the real computing is performed. The majority of the nodes in a cluster are typically compute nodes. Based on the scheduling of the cluster, the compute node that is given one or more tasks to execute is part of the overall solution.

Note: The standard compute node in the IBM Cluster 1350 is the xSeries Model 335. It is a 1U server with one or two Intel Pentium 4 (Xeon) processors. The system also has a 133 MHz front-side bus, integrated Ultra160 SCSI interface, two integrated 10/100 Ethernet NICs, and two 64-bit/33 MHz PCI slots (one full-length and one half-length).

IBM BladeCenter can also be used as a compute node in the IBM Cluster 1350. While the BladeCenter is a 7U chassis, it contains up to 14 blades. Because each blade is a separate system with dual processors, disk storage, and networking, the BladeCenter provides the greatest processor and storage density of any cluster node option.

Infrastructure node

Infrastructure nodes perform secondary functions (the primary function of most clusters is to compute as fast and as reliably as possible) which are still essential for the cluster to operate. These functions include:

- ▶ Control node: Control nodes provide services that help other nodes in the cluster work together to obtain the desired result. Control nodes can provide two sets of functions:
 - Infrastructure services like DHCP, DNS, and other similar functions for the cluster: These functions enable the nodes to easily be added to the cluster and to ensure that they can communicate with the other nodes.
 - Scheduling what tasks are to be done by what compute nodes: For example, if a compute node finishes one task and is available to do additional work, the control node may assign that node the next task that requires the work.

Note: Either an xSeries 300 Model 335 or 345 can be used as a cluster control node.

- ▶ Storage node: For most applications that are run in a cluster, compute nodes must have fast, reliable, and simultaneous access to the storage system. This

can be accomplished in a variety of ways depending on the specific requirements of the application. In some cases, storage may be direct-attached to compute nodes. In other cases, the storage requirements may be filled through one or more network shared devices. In this case, storage nodes enable and control access to the storage subsystem.

Note: The storage node of Cluster 1350 is the xSeries Model 360. The Model 360 is a 3U server with up to four processors. While the x360 can host internal fibre channel disks, it can also connect to other SAN and NAS devices.

- ▶ **Install node:** In most clusters, the compute nodes (and other nodes) may need to be reconfigured or reinstalled with a new image relatively often. The installation node provides the images and the mechanism to easily and quickly install or reinstall software on the cluster nodes.
- ▶ **Management node:** Clusters are complex environments. Therefore, the management of the individual components is very important. The management node provides many capabilities, including:
 - Monitoring the status of individual nodes
 - Issuing management commands to individual nodes to correct problems, or to provide simple (yet useful) commands to perform such functions as power on/off
 - Handling events or alarms originating from the various nodes in the cluster

Note: The Model 345 is used as the management node (although it can also be used as a storage or computer node), and is configured as a 2U unit with two Pentium 4 (Xeon) processors. The x345 has an Advanced Systems Management processor (also known as a *service processor*), which allows remote node power on/off/reset capability, monitors node environmental conditions (fan speed, temperature, power, etc.), and allows remote Power-On Self-Test/Basic Input Output System (POST/BIOS) console, power management, and SNMP alerts.

- ▶ **User node:** The individual nodes of a cluster are often on a private network that you cannot access directly from the outside or corporate network. Even if they are accessible, most cluster nodes are not necessarily configured to provide an optimal user interface. The user node is the one type of node that is configured to provide that interface for users (possibly on outside networks), who may gain access to the cluster to request that a job be run, or to access the results of a previously run job.

Node types

Now that you understand the types of logical functions that typically exist in a cluster, we discuss the node types (from a hardware perspective) that may exist in a cluster. Generally, nodes within a cluster are categorized as one of three types, each of which is explained in the following sections.

Head nodes

Head node is a generic term. The head node aids in controlling the cluster, but can be used in additional ways. Head nodes generally provide one or more of the following logical node functions:

- ▶ User node
- ▶ Control node
- ▶ Management node
- ▶ Installation node

In a small cluster, typically no more than 16 nodes, all of these functions can be combined in one head node. In larger clusters, the functions are best split across multiple machines for security and performance reasons.

Compute nodes

The compute nodes form the heart of the cluster. The user, control, management, and storage nodes are all designed to support the compute nodes.

Most computations are actually performed on the compute nodes. These nodes are logically grouped, depending on the needs of the job and as defined by the job scheduler.

Storage nodes

Often when discussing cluster structures, a storage node is defined as a third type of node. However, in practice, a storage node is often a specialized version of either a head node or a compute node.

The reason that storage nodes are sometimes designated as a unique node type is that the hardware and software requirements to support storage devices might vary from other head or compute nodes. Depending on your storage requirements and the type of storage access you require, this may include special adapters and drivers to support RAID-5, storage devices attached through channel storage, and others.

10.2.2 Networks

Nodes within a cluster have many reasons to communicate. After all, they are cooperating to provide an overall result. The needs defined by the different

communication types tend to dictate the various network technologies that work to serve those needs.

Communication types

Some of the key communications that must occur are explained in the following sections.

Interprocess communication

The nodes in a cluster work together to provide an overall solution. Therefore, the task that each node performs is often interrelated with the tasks that the other nodes perform.

Just as in a multi-processing or multi-threaded single system, there is almost always a need for communication between processes to coordinate processing and handle concurrent access to shared resources. The interprocess communication network is often referred to as the *cluster interconnect*. It may share a physical connection with the storage or management networks.

Management operations

With a large number of nodes, it is important to recognize when an individual node may be failing. It is also important to proactively issue commands that may prevent an error from occurring.

It is not feasible to manage multiple nodes of a cluster individually in a stand-alone manner. There must be a way to communicate from one or more management nodes to any other node in the cluster.

Note: IBM Cluster 1350 provides a service processor network. It uses the newer RSA adaptors, which are one for every 11 nodes. The service processor is built in and has its own Ethernet port and cable. Because of this, only one Ethernet cable per eight nodes is required, as opposed to using a PCI slot and Ethernet cable per node.

Software installation

For a variety of reasons, individual nodes often need to have their software reinstalled. This may be due to reconfiguration for a specific task to be performed by the cluster, or simply to replace the operating environment when an individual node has suffered a disk failure.

Either way it is important that system images can be efficiently defined and distributed. Since this operation may be done to certain nodes while the rest of the cluster is operational and performing work, it is important that pushing a new operating environment to some nodes does not adversely affect the performance of the other nodes.

Storage access

In most clusters, individual nodes need to read or write data to a shared storage device. For high performance computing, it is contrary to the intention of the cluster if this storage access proved to be a bottleneck and resulted in other nodes waiting to access the device. Therefore, it is important to have both an efficient network for fast access to storage, and a file system that provides efficient, concurrent access to the data.

Common network technologies

Depending on the actual application and the performance requirements, these various communications should be carried out often on different networks. Therefore, you typically have more than one network and network type linking the nodes in your cluster. Common network technologies used in clusters are explained in the following sections.

Fast Ethernet

Fast Ethernet and TCP/IP are the two standards used for most networking. They are simple and cost-effective, and the technology is always improving. Fast Ethernet works well with lower speed Intel-based machines used for Beowulf nodes, and can operate at either half-duplex or full-duplex. With full-duplex, data can be sent and received at the same time. Full-duplex transmission is deployed either between the ports on two switches, between a computer and a switch port, or between two computers. Full-duplex requires a switch, since a hub will not work.

Gigabit Ethernet

Gigabit Ethernet uses a modified version of the American National Standards Institute (ANSI) X3T11 Fibre Channel standard physical layer (FC-0) to achieve one Gigabit per second raw bandwidth. Gigabit Ethernet supports multi and single-mode optical-fiber and short-haul copper cabling. Fibre is ideal for connectivity between switches and servers and can reach a greater distance (from 500 meters to 10 kilometers today) than copper.

Myrinet™

Myrinet is a high-performance, high-availability, packet communication and switching technology. Myrinet links may be up to 200m in length on multimode fiber, and operate full-duplex at a data rate of 2+2 Gigabits per second. Myrinet switches are simple, low-latency, cut-through devices that can be combined to provide networks up to thousands of hosts and network-bisection data rates in Terabits per second. Myrinet interfaces include a processor and firmware to offload packet processing from the host computer.

Myrinet software support for Linux and for other operating systems provides multiple software interfaces. This includes the native Myrinet API, TCP/IP and UDP/IP, MPI, VI, PVM, and Sockets. All Myrinet software support is open source.

For more information about Myrinet, see the following Web site:

<http://www.myri.com/>

Fibre Channel

Fibre Channel is a data storage technology that can provide cluster nodes with shared access to data. Fibre Channel is a loop architecture as opposed to a bus architecture like SCSI. The loop structure enables the rapid exchange of data from device to device, achieving a maximum data transfer rate of 100 MB/sec. Fibre Channel is a common technology used to provide clusters with rapid and shared data access.

Note: Cluster 1350 supports the IBM FASTT Fibre storage subsystem, and the IBM TotalStorage Enterprise Storage Server.

Internet Small Computer Systems Interface (iSCSI)

iSCSI is a standard that allows SCSI information to be encapsulated and transported over a TCP/IP network. The benefit of this approach is that organizations can take advantage of their existing network infrastructure to provide access to data storage resources anywhere on the network using the same I/O commands used for local data. iSCSI can allow Linux clusters to access remote data stores without requiring additional infrastructure.

As the iSCSI standard matures, several developments are expected. The first is the introduction of Host Bus Adapters (HBAs). HBAs improve performance and throughput by offloading the TCP/IP overhead from client operating systems. The second is the deployment of 10-Gigabit Ethernet, which dramatically improves the data rates at which iSCSI can operate.

Cluster 1350 does not support iSCSI. It is mentioned here for completeness.

Note: IBM TotalStorage IP Storage 200i is an iSCSI solution that provides support for 109.2 GB up to 3.52 TB of RAID storage.

InfiniBand

InfiniBand is a technology that is designed to provide a channel-based switched fabric that is capable of handling any kind of communication between processors and I/O devices. Because of the high bandwidth designed into the standard (2.5 Gbits, 10 Gbits, and 30 Gbits/sec.), InfiniBand is expected to gradually replace the existing Peripheral Component Interconnect (PCI) shared-bus

approach used in most of modern Intel-based computers, and is seen as an ideal communication media for the data-intensive demands of clusters.

Cluster 1350 does not support InfiniBand. It is mentioned here for completeness.

Note: IBM offers several InfiniBand products including the InfiniBlue 4X Host Channel Adapter (HCA). This allows current PCI-based Linux servers to connect to an InfiniBand network.

10.2.3 Support software

Software binds the hardware infrastructure together. It creates and manages the abstraction layer that allows users and applications to view the collective cluster resources as a single computing entity. There are three main cluster software categories.

Operating system

All nodes in a cluster typically run the same operating system. For large clusters, it is important that this operating system be economically viable, as well as easily deployable across the nodes. This implies that it should be easily installable in an automated fashion, and support remote access for management and configuration.

Linux meets these requirements. In fact, the low cost of obtaining Linux and the rich function for developing and deploying applications that run on Linux have been the catalyst for increased interest in using clusters to meet the needs of a variety of computing problems. By using relatively low cost Intel-based hardware and Linux, many businesses that previously could not justify the expense of a cluster are now adopting the technology at a frenetic pace. As we previously mentioned, there are other considerations when building a cluster such as how to manage a large number of nodes.

Note: Currently, the only supported Linux distributions for Cluster 1350 are provided by Red Hat and SuSE. Additional Linux distributions and versions will be considered for the future.

File system

Depending on the actual application, many clusters require a large number of nodes to have concurrent and very high-speed access to the same files. In these cases, it is important to provide a shared file system that provides data protection and the performance that is typically required for cluster-based applications.

Note: IBM provides the General Parallel File System (GPFS) for Linux. GPFS is a proven file system that has been used for years in AIX clustering environments. It is available for Linux and xSeries hardware.

Management tools

When given the hardware previously discussed, you can fairly easily construct a cluster. This is only the beginning of putting a cluster into production.

One of the most difficult parts of deploying a cluster is having the right tools in place to manage a large number of nodes as a single entity. This includes:

- ▶ Managing node and node group information
- ▶ Monitoring and controlling hardware
- ▶ Running remote commands on multiple machines
- ▶ Monitoring system events
- ▶ Managing and synchronizing configuration files for all machines
- ▶ Providing security
- ▶ Diagnosing problems
- ▶ Installing software

Note: IBM provides cluster management on Cluster 1350 through two utilities: xCAT and Cluster Systems Management (CSM) for Linux. xCAT is a more mature product and uses the xSeries service processor. CSM is a newer tool and includes technology derived from PSSP. Integration between CSM and xCAT is provided by an add-on to CSM called Enhanced Cluster Tools. For more information on ECT, see:

<http://www.alphaworks.ibm.com/tech/ect4linux>

10.2.4 Applications

Few off-the-shelf HPC applications exist for the simple reason that HP clusters tend to be focused on a specific, unique task. HTC solutions also tend to be custom-written and focus on massive calculation projects. HA clusters typically rely on middleware as the software layer that requires the most customization to ensure reliability, recoverability, and even scalability.

The two HA subtypes (storage and database clusters) each have their own special concerns that may require application customization. For storage clusters, more emphasis is placed on the capabilities of the distributed file system itself. However, some customization of specialized retrieval libraries (layered between the application and the distributed file system) is common.

With database clusters, customization is driven by the primary concerns of availability and recovery. However, scalability is increasingly becoming a concern.

You can approach the load balancing clusters in two ways:

► **Server farm**

The traditional way of scaling loads is to provide a server farm that sits behind a server that spreads all incoming requests across the various nodes of the cluster. The advantage to server farms is the low amount of customization typically required to implement this type of load balancing.

► **Process migration/Single System Image (SSI) cluster**

Using kernel-level extensions, process migration clusters can seamlessly move execution processes to and from any node in the cluster. The benefit of this approach is that applications do not have to be rewritten to take advantage of this approach to load balancing.

An SSI cluster expands the simple process migration model to create unified cluster-wide inter-process communication (IPC) objects such as pipes, fifos, semaphores, message queues, shared memory, sockets, and signals. This allows programs, processes, and users to work together regardless of the node in the cluster on they are executing.

Either type of load balancing (server farm and process migration) may still require a distributed file system or additional customization to deal with data concurrency issues.

Figure 10-8 shows how all the various cluster components come together. Notice that the top four layers are all software components. This abstraction layer serves to allow the creation of a cluster environment from a heterogeneous mix of hardware. Because an effective cluster can be made up of any type or brand of server, interconnect, or storage subsystem, clusters can take advantage of commodity hardware or the existing IT resources of an organization.




Figure 10-8 Layout of cluster components

Note: IBM provides the following cluster-enabled applications:

- ▶ DB2 Universal Database (UDB) Enterprise-Extended Edition (EEE)
- ▶ WebSphere Application Server
- ▶ WebSphere Commerce Suite
- ▶ WebSphere Edge Server
- ▶ Telecom Application Server
- ▶ WebSphere Everyplace Suite

10.3 Linux cluster history

The current state of Linux cluster solutions is best understood through its history. Modern Linux clusters are conceptually indebted to open source software, the Intel processor, and the demands of the HPC community. There were clusters of non-Intel/non-Linux machines before (and after) the now-famous Beowulf cluster program started in 1994 at NASA. However, the idea of a cluster was not truly feasible or economical until the commodity nature of PC-based hardware was combined with the open source nature of the Linux operating system.

The Beowulf program

The Beowulf program started with a simple idea: Use 16 inexpensive Intel-based PCs and open source parallel programming libraries to create a virtual supercomputer environment. Fast Ethernet was too slow to keep up with the bus speed of the individual nodes (bus speed continues to be a major roadblock even in today's clusters). Because of this, Donald Becker (a NASA contractor) rewrote a network card driver to support *channel bonding*. Channel bonding is a process that combines the four network cards in each into a single, virtual connection that is fast enough to permit parallel processing software running on each computer to communicate as if they were part of a single system.



Figure 10-9 Wiglaf, the first cluster in the Beowulf program

All the software that made up the system was open source, including the Linux kernel, the NIC driver, and the parallel processing libraries. Considering this, the NASA team (Becker, Tom Sterling, John Salmon, and other CESDIS members) continued and developed the ability to tune the cluster to the specific characteristics of the system.

Finally, because all of the hardware involved was standardized and inexpensive, additional nodes could be added (or replaced) as necessary. Today, the phrase *Beowulf cluster* refers to a class of cluster where:

- ▶ The majority of the components are commodity off-the-shelf (COTS).
- ▶ The operating system is open source (Linux or BSD variants).
- ▶ There is emphasis of the cluster on HPC and as an Industry Standard Message Passing interface.

Emergence of high availability clusters

Starting in the 1990s, Linux was increasingly used in an infrastructure role providing basic (and essential) services including routing, firewall/IDS, Web, and DHCP/DNS. The need for a higher level of availability also increased along with the penetration of Linux into businesses and government.

The Linux high availability project started with the purpose of bringing failover and takeover capabilities to the Linux operating system. The central component of the Linux high availability project is the *heartbeat program*. The heartbeat program works on a two-node model: two identically-configured computers, with one designated as the “primary”, and one as the “secondary”. The secondary node monitors the “heartbeat” or continual signal sent by the primary through a network connection or serial cable. When the primary node fails (and fails to

send the heartbeat signal), the secondary activates its services and takes ownership of primary's IP address, maintaining the cluster service.

While Linux high availability provides a single backup for essential infrastructure servers, the two-node limitation limited its effectiveness. As a result, administrators increasingly began to combine the heartbeat capabilities of Linux high availability with the availability and load balancing capabilities of another project, the *Linux Virtual Server* (LVS). LVS works by assigning the public IP address of the cluster to a load-balancing server. The load balancing server then retransmits requests for services to the least-busy node of the cluster. When the load balancing LVS is mirrored with Linux high availability, the result is a highly available and scalable cluster for Web-based services.

Linux load balancing

In the mid 1990s, Professor Amnon Barak of the Hebrew University in Jerusalem ported the code (originally developed for BSD UNIX) of the Multicomputer Operating System for High Performance Cluster Computing (MOSIX) to Linux. MOSIX uses special algorithms to automatically (and optimally) move processes around the cluster. A process can be started from any node, and can be moved to or from any other node. Because MOSIX is implemented as a kernel extension, Linux applications do not have to be rewritten to take advantage of its load balancing capabilities.

In 2002, MOSIX began offering an optional parallel I/O (MOPI) component. This component attempts to improve I/O operations by migrating parallel processes to the nodes that hold that data to be used by the processes.

In 1998, some members of the original Beowulf program formed the Scyld Computing Corporation. Scyld Beowulf added extensive cluster installation and management software to the collection of open source cluster tools then in existence. It also incorporated the idea of process migration. Similar to MOSIX, Scyld Beowulf uses an extension to the Linux kernel called Bproc to migrate cluster processes from the master node to other nodes in the cluster. In addition, all process management can be performed from the master node. As a result, individual cluster nodes are not required to have duplicate copies of cluster application code. This simplifies cluster administration and maintenance and frees node storage space for application data and cache.

There is a major difference between MOSIX and Scyld. MOSIX (and openMosix) transparently migrates processes under kernel control based on a least loaded algorithm, with open file handles. Scyld migrates processes under user control, closing file handles. This makes the process faster and leaner to migrate, but necessitates reopening files that you want to access.

In late 2001, Professor Barak decided to discontinue the open source license of MOSIX. Moshe Bar, a former developer of the Linux MOSIX code decided to continue open source developments under the name *openMosix*. openMosix has incorporated several features that are not found in the original MOSIX, including support for user-mode Linux (UML), among others. It continues to be one of the most popular Linux cluster software components.

Note: UML is part of the Linux kernel that allows multiple virtual machines, each running their own version of the Linux kernel, to run on top of a single host kernel. Because the host kernel provides each virtual machine with simulated hardware, they can duplicate nearly all programs that are available on the host itself. UML is used by developers to test and debug Linux kernels and applications in a secure environment.

Management consolidation

While various open source cluster software components were maturing in the 1990s, management tools still lacked maturity and integration. In April of 2000, the Open Cluster Group, a collection of companies (IBM, Intel, Dell, SGI, Veridian, MSC.Software Corporation) and research organizations (Oak Ridge National Laboratory (ORNL), Indiana University, University of Illinois National Center for Super Computing Applications (NCSA)) released their first project: the *Open Source Cluster Application Resources* (OSCAR).

The goal of the OSCAR project is to identify the best open source utilities for building, programming, and administering HPC clusters. Once identified, these programs are bundled together and distributed as a single package. OSCAR has been an extremely popular open source cluster package, and is used by a wide range of organizations.

Note: IBM provides tools for CSM to work with OSCAR through its alphaWorks site. For more information, see:

<http://www.alphaworks.ibm.com/tech/ect4linux>

Another approach to cluster management and infrastructure is the *Single Systems Image project*. SSI attempts to provide a general-purpose Linux clustering environment using the best open source components for availability, scalability, and manageability. SSI (and the Cluster Infrastructure (CI) project it relies on) has its origin in Hewlett Packard NonStop Clusters for UnixWare technology. The CI code is responsible for maintaining membership of and communication between the various cluster nodes. SSI creates the appearance of a single system through shared process, IPC, device, and networking space, a clustered file system (through OpenGFS), load leveling (using the MOSIX algorithm), and a single management interface.

Note: SSI was voted the best open source project of the 2002 LinuxWorld Conference and Expo.

10.4 Linux cluster solutions

IBM offers two main Linux cluster solutions: Cluster 1350 based on xSeries, and Cluster 1600 based on pSeries.

10.4.1 Cluster 1350

Cluster 1350 is designed as a solution for industrial, financial services, life sciences, governmental, and educational organizations. Because Cluster 1350 is based on the xSeries, the customer is provided with industry-standard equipment along with an extensive capability to customize the cluster capabilities.

Each node (or, in the case of BladeCenter, each blade) is dual-processor capable, and the client can choose between Intel Xeon or AMD Opteron processors, storage options that range from internal disks to external NAS or SAN technologies, and cluster VLAN (also referred to as the cluster interconnect network) options that include Gigabit Ethernet and even Myrinet-2000. For scalability, Cluster 1350 is provided in a default configuration of a single management node and up to 511 computer nodes, 32 of which can provide specialized cluster file sharing. Larger clusters or more specialized configurations are available, and are handled through a special order process.

Because cluster software is as important as the cluster hardware, IBM provides several tools to help manage and support cluster operations, which are featured in the following sections.

Cluster Systems Management (CSM)

CSM provides for cluster installation, management, and maintenance capabilities. In particular, CSM allows for centralized node configuration, cluster-wide event management (such as process failures), and remote



hardware control (including the ability power off and on individual nodes) all through a single interface.

Cluster Systems Management (CSM) for Linux is designed for simple, low-cost management of distributed and clustered IBM xSeries servers in technical and commercial computing environments. CSM, included with the IBM Cluster 1350, dramatically simplifies administration of a cluster by providing management from a single point-of-control.

Note: Because CSM supports both Cluster 1350 and Cluster 1600, it can serve as a central management utility not only for separate clusters, but can also be used to create a hybrid cluster of both xSeries and pSeries cluster nodes.

Enhanced Cluster Tools for Linux

Enhanced Cluster Tools (ECT) for Linux is a set of tools that complement CSM and enhance the management of Linux clusters. ECT for Linux adds to CSM features like remote access to hardware inventory and vitals, and remote access to service processor logs.

For more informations, see:

<http://www.alphaworks.ibm.com/tech/ect4linux>

General Parallel File System (GPFS)

GPFS is used to provide shared-disk access to the entire cluster, while still retaining scalability and high performance. Additionally, GPS allows NFS-capable clients outside the cluster to access cluster file resources. Because GPFS uses standard UNIX file system interfaces, parallel applications running across multiple nodes of the cluster as well as serial applications running on a single node can readily access shared files. Finally, for high availability concerns, GPFS can be configured for failover from both disk and server malfunctions.

Note: For more information on Cluster 1350, see:

<http://www.ibm.com/servers/eserver/clusters/hardware/1350.pdf>

Higher performance and scalability

GPFS is particularly appropriate in an environment where the aggregate peak need for data exceeds the capacity of other shared file systems. To support its performance objectives, GPFS is implemented using data striping across multiple disks and multiple nodes, and it employs client-side data caching. GPFS

provides large blocksize options for highly efficient I/O and has the ability to perform read-ahead and write-behind file functions.

GPFS uses block level locking based on a sophisticated token management system designed to provide data consistency while allowing multiple application nodes concurrent access to a file. When hardware resource demands are high, GPFS can find an available path to the data by using multiple, independent paths to the same file data from anywhere in the cluster.

Scalability is provided by increasing the number of nodes and disks assigned to support the I/O as the overall cluster configuration grows. A scalable interconnect such as the Myricom Myrinet™-2000 is generally used to support this overall scaling.

High-performance disk configuration

There are two configuration options that may be used with GPFS for Linux: the Network Shared Disk Server Model and the Directly Attached Model.

The Network Shared Disk (NSD) Server Model is analogous to the Virtual Shared Disk (VSD) implementation used by GPFS for AIX 5L. A subset of the total node population is defined as NSD storage nodes, and GPFS disk drives are attached only to these nodes. NSD implements a software layer that ships I/O requests from an application node to the cluster interconnect fabric, which then passes that request to an NSD storage node to perform the I/O, and passes data back to the application node. This remote I/O is transparent to the application requesting it. For storage nodes, RAID disk configurations using the IBM FASTT disk are supported, and disk configurations should be designed with alternate node and adapter paths to permit failover. A scalable interconnect fabric is recommended for the NSD model.

In the Directly Attached Model, each node that mounts a GPFS file system must have a direct connection to all of the disks in the GPFS file system. This disk attachment model requires the use of a Fibre Channel Switch and a Storage Area Network (SAN) disk configuration. In supported environments, each node that mounts GPFS must have a FASTT Host Adapter installed for connection to the Fibre Channel switch.

Availability and recoverability

GPFS for Linux is designed to survive many system and I/O failures. Through use of the IBM Reliable Scalable Cluster Technology (RSCT) included with GPFS, it is designed to automatically recover from node, disk connection, and disk adapter failures.

RSCT includes the following three subsystems that collectively provide a high availability infrastructure to IBM clusters:

- ▶ *System Resource Controller* provides a common way to start, stop, and collect status information on processes and subsystems.
- ▶ *Group Services* provides distributed coordination and synchronization service, and monitors the status of all processes in a group informing members upon failure.
- ▶ *Topology Services* provides adapter status and node connectivity information using a heartbeat. It also provides a reliable messaging service.

GPFS is designed to transparently failover lock servers and other GPFS central services based on the monitoring and coordination facilities of RSCT. Similarly, GPFS can continue to operate in the event of disk connection failures when disks are connected to two servers.

Data replication is supported by GPFS and may be used to help further reduce the chances of losing data if storage media fail, or if RAID storage is not implemented. GPFS is a logging file system that allows the recreation of consistent structures for rapid recovery after node failures.

10.4.2 Cluster 1600

Cluster 1600 is a proven pSeries-based cluster offering that now supports Linux. Cluster 1600 will commonly be used in the following scenarios:

- ▶ Server consolidation: Lower costs can be realized by consolidating applications and streamlining data center operations.
- ▶ Large-scale databases: Cluster 1600 provides scalability with excellent price/performance for the multi-terabyte enterprise data warehouses.
- ▶ High performance: Fast pSeries servers are ideal for large-scale computational models in science and engineering.



Cluster 1600 is based on SP and pSeries technology, which means that clients can mix and match servers to fit their existing infrastructure and needs. Additionally, Cluster 1600 is scalable up to a total count of 128 nodes (although a maximum of 64 physical systems), which can consist of pSeries, SP, or AIX/Linux operating system images (LPARs).

Like the Cluster 1350, Cluster 1600 provides CSM for management, GPFS for file storage, and can support Myrinet-2000 for the cluster VLAN. In addition, IBM makes the following products available for Cluster 1600:

- ▶ Engineering and Scientific Subroutine Library (ESSL) for pSeries Linux 4.1: This library contains a state-of-the-art collection of mathematical subroutines that provide optimum performance for floating-point engineering and scientific

applications. ESSL contains over 400 high-performance mathematical subroutines.

- Parallel ESSL for pSeries Linux 3.1: This add-on to ESSL contains over 100 high-performance mathematical subroutines specifically designed to exploit the full parallel power of a pSeries/Linux cluster.

Note: For more information on Cluster 1600, see:

http://www-1.ibm.com/servers/eserver/clusters/hardware/1600_facts.pdf

10.4.3 ClusterProven®

The IBM ClusterProven program sets specific criteria for validating end-to-end solutions that meet industry standards for high availability on every IBM platform. Developers who achieve ClusterProven validation earn the use of the ClusterProven mark. This allows them to use ClusterProven in marketing their solution to customers and receive co-marketing assistance from IBM.

To see a list of ClusterProven applications, go to:

<http://www-1.ibm.com/servers/clusters/clusterproven.pdf>

10.4.4 Solution examples

Cluster 1350 and Cluster 1600 are the primary current Linux cluster packages offered by IBM. However, companies still choose IBM hardware and middleware to support their custom cluster solutions. This section offers examples of Linux clusters in the real world.

High-performance computing: Blue Gene/L

Blue Gene is an IBM supercomputing project that aims to build a new family of supercomputer optimized for bandwidth, scalability, and the ability to handle large amounts of data while consuming a fraction of the power and floor space required by today's fastest systems. Built for the Lawrence Livermore National Laboratory in California, USA, by IBM, Blue Gene has already made history. On November 18th, 2003, the Blue Gene/L prototype ranked as the 73rd fastest computer systems in the world, according to the Top500 Supercomputer project.




Note: For the actual listing of the November, 2003 edition of the Top500 Supercomputer project see:

<http://www.top500.org/list/2003/11/>

The full Blue Gene/L machine, which is being built for the Lawrence Livermore National Laboratory in California, will be 128 times larger than the prototype, occupying 64 full racks (the prototype was just one rack). When completed in 2005, Blue Gene/L is expected to lead the Top500 supercomputer list. Compared with today's fastest supercomputers, it will be six times faster, consume 1/15th the power per computation, and will be 10 times more compact than today's fastest supercomputers.

Blue Gene/L overview



Blue Gene/L is a scalable system with the maximum size of 65,536 compute nodes. Each node is implemented on a single CMOS chip (which contains two processors, caches, 4 MB of EDRAM, and multiple network interfaces) plus external memory. The systems-level design puts two nodes per card, 16 cards per node board, and eight node boards per 512-node midplane. Two midplanes fit into a rack. Each processor is capable of four floating point operations per cycle (such as two multiply-adds per cycle). This yields 2.8 Tflop/s peak performance per rack. A complete BG/L system is 64 racks.

Blue Gene/L network design

Blue Gene/L has five different networks that accomplish the various calculation, I/O, and management functions of the cluster:

- ▶ Three Dimensional Torus: The main communication backbone of Blue Gene/L, this network enables point-to-point messaging between compute nodes. This design allows for a total torus interconnect bandwidth approaching 134 TB/s.
- ▶ Global combining tree: This network supports a global combining/broadcast tree for collective operations. With approximately 1.4 GB/s of bandwidth between any given node pair in the cluster, a signal can traverse the tree in less than one microsecond.
- ▶ Global barrier tree: This network supports global interrupt actions across the cluster.
- ▶ Remote management: This Gigabit Ethernet network supports functions such as booting, diagnostics, and host control.
- ▶ I/O: Also implemented using Ethernet, this network is used for data transfers and host communications. For cost and overall system efficiency, compute nodes are not hooked directly up to the Gigabit Ethernet, but rather use the global tree for communicating with their I/O nodes.

Why Blue Gene?

While the original and overriding goal of the Blue Gene project is to build a petaflop scale machine to perform calculations in the area of life sciences, an additional goal is to attempt to provide insight into the following areas:

- ▶ Simulation methodologies, analysis tools, and biological systems
- ▶ Scientific and technical computing applications in other domains such as climate and materials science
- ▶ Autonomic and on-demand technologies in a high performance environment

High-performance computing: NCSA

The National Center for Supercomputing Applications (NCSA) is a research organization based in the United States. Its mission is to create high-performance computing infrastructures for scientists. IBM has installed two Linux clusters for the NCSA. The first cluster includes more than 600 xSeries running Linux and Myricom's Myrinet cluster interconnect network. The second cluster is one of the first to use Intel's next generation 64-bit Itanium processor and will run SuSE.



Figure 10-10 The NCSA Linux cluster built on the xSeries server

It would take more than one-and-a-half million years for a person to tabulate with a calculator the number of calculations that the NSCA supercomputer handles in a single second. The NSCA needs immense speed and power to simulate the violent collision of black holes and the resulting gravitational waves.

You can learn more about the IBM NCSA cluster on the Web at:

<http://www.ncsa.uiuc.edu/About/NCSA/>

Plus, you can find more information about Linux clusters on the NCSA Web site at:

<http://www.ncsa.uiuc.edu/News/Access/Stories/ClusterUpdate/>

Lawrence Livermore National Laboratory in 1998. That simulation was done over one week using a 3,840-processor computer. However, follow-up work was never done because the Livermore machine was relegated to a classified environment.

HA Clusters Computing: BBDO Dusseldorf

BBDO Dusseldorf is a German-based company that provides complete Web hosting and infrastructure solutions. To provide Web-hosting for its clients, BBDO INTERACTIVE has chosen an IBM solution including DB2 Universal Database and xSeries servers running the open source Linux operating system. BBDO Dusseldorf maintains high availability of client Web sites by using LifeKeeper™ for DB2 clustering software from IBM Business Partner SteelEye.




Figure 10-11 BBDO Dusseldorf high availability cluster on the xSeries server

You can read more about the BBDO Dusseldorf cluster at:

<http://www.eweek.com/article2/0,3959,131649,00.asp>

Grid computing: National Digital Mammography Archive (NDMA)

The NDMA is a distributed computing infrastructure that connects hospitals at the University of Chicago, the University of North Carolina, and the Sunnybrook and Women's College Hospital in Toronto. The purpose of the NDMA is to deliver computing resources as a utility-like service over a secure Internet connection. This enables up to thousands of hospitals to store mammograms in digital form, giving medical personnel near-instantaneous access to patient records.

The NDMA portal systems run almost exclusively on IBM hardware including 16 distributed Netfinity servers running Linux. You can learn more about the NDMA cluster on the Web at:

<http://nscp.upenn.edu/>

10.5 The future of Linux clusters

Clusters are now used as the building block of other emerging IT models and technologies. In this way, clusters are becoming components for larger projects, as opposed to the traditional view of a cluster as a stand-alone system that performed a specific computing role.

10.5.1 Emerging trends

There are three major emerging trends that rely extensively on the features provide by Linux. These trends are explained in the following sections.

Grid computing

Grid computing enables the virtualization of distributed computing and data resources such as processing, network bandwidth, and storage capacity to create a single system image, granting users and applications seamless access to vast IT capabilities. Just as an Internet user views a unified instance of content through the Web, a grid user essentially sees a single, large virtual computer.

At its core, grid computing is based on an open set of standards and protocols such as Open Grid Services Architecture (OGSA), which enable communication across heterogeneous, geographically dispersed environments. With grid computing, organizations can optimize computing and data resources, pool them for large capacity workloads, share them across networks, and enable collaboration.

In this way, a grid can be seen as a *metaccluster*. A metaccluster is a collection of clusters that are interconnected with wide area network (WAN) technologies to

provide seamless access to the combined resources of the grid. Grids are distinguished from conventional distributed computing by their focus on large-scale resource sharing, innovative applications, and in some cases, high-performance orientation.

IBM has a long and thorough involvement with both the technology and the business issues that have led to the grid computing evolution. *Virtualization* — the driving force behind grid computing — has been a key factor since the earliest days of electronic business computing.

Grid computing goes far beyond sheer computing power. Today's operating environments must be resilient, flexible, and integrated as never before. Organizations around the world are experiencing substantial benefits by implementing grids in critical business processes to achieve both business and technology benefits.

With grid computing you can:

- ▶ Accelerate time to result
- ▶ Enable collaboration and promote operational flexibility
- ▶ Efficiently scale to meet variable business demands
- ▶ Increase productivity
- ▶ Leverage existing capital investment

You can learn more about grid computing on the Web at:

<http://www.ibm.com/grid/index.shtml>

IBM Grid Toolbox for Linux

The IBM Grid Toolbox is a set of installable packages that includes the Globus Toolkit 2.2 with additional documentation and custom installation scripts written for IBM eServer hardware running Linux.

The IBM Grid Toolbox is an integrated set of tools and software that facilitate the creation of grids and applications, which can exploit the advanced capabilities of the grid using a combination of this toolbox and other technologies.

The components of the IBM Grid Toolbox can be used either independently or together to develop useful grid applications and programming tools. The IBM Grid Toolbox components are as follows:

- ▶ Globus Resource Allocation Manager (GRAM)
- ▶ Grid Security Infrastructure (GSI)
- ▶ Monitoring and Discovery Service (MDS)
- ▶ Global Access to Secondary Storage (GASS)
- ▶ Globus Replica Catalog
- ▶ Grid Toolbox I/O
- ▶ Simple CA

For more information about IBM Grid Toolbox, see:

<http://www.alphaworks.ibm.com/tech/gridtoolbox>

IBM Grid Value At Work

IBM Grid Value at Work assists in developing a grid deployment roadmap. The roadmap includes components for total cost of ownership (TCO) analysis, grid capacity planning, and business value estimation.

IBM Grid Value at Work gives executives the flexibility to focus on one or more of these components to determine the appropriate investment levels for grid computing. It also helps develop financial models including return on investment (ROI) and internal rate of return (IRR), as well as dollar value outputs such as net present value (NPV), cash flow and payback period.

Additionally, the IBM Grid Value at Work technique helps identify a grid's most appropriate usages by analyzing the comprehensive business value of the implementation across all deployment stages, and quantifying the following information:

- ▶ **Business analytics:** The business value of improving the speed, accuracy, frequency, or throughput of the analysis
- ▶ **Research and development/engineering and design:** IT infrastructure cost savings and potential benefits from faster time to market or improved design quality
- ▶ **Enterprise optimization:** Cost savings from virtualization of the IT infrastructure, and the value derived from exploiting existing, under-utilized resources
- ▶ **Government:** Cost savings from integrated information access and shared usage of IT resources across departments and organizations

You can download the complete details, at:

http://www.ibm.com/grid/pdf/grid_value_at_work.pdf

Autonomic computing

Autonomic computing is an attempt to manage the complexity of IT systems by using the inspiration found in one of the most complex systems of the human body, the *autonomic nervous system*. The focus of autonomic computing is to design and build computing systems capable of running themselves. It requires adjusting to varying circumstances, and preparing their resources to handle most efficiently the workloads placed upon them.

To read more about autonomic computing, see:

<http://researchweb.watson.ibm.com/autonomic/>

On demand computing

On demand computing describes an environment where users own IT assets, but can dynamically increase and decrease them based on need. The main advantage for users is variable pricing based on usage, similar to the way energy utility companies provide and bill for their services. In addition, the nature of on-demand computing allows users to outsource the management of such a demand-driven system to third parties.

You can read more about on demand computing on the Web at:

http://www.ibm.com/services/ondemand/index_flash.html

10.5.2 The role of Linux clusters in emerging trends

Linux clusters will serve grid computing by supporting the intensive computational and storage demands imposed by the distributive nature of the grid. In addition, Linux clusters provide the least expensive way of providing the fault tolerance demands of any autonomic system. Linux clusters are the most cost-effective solution for providing the scalability requirements of on-demand computing.

Another future view of Linux clustering is as an *HPC appliance*. The HPC appliance will be based on super thin nodes, with little or no persistent storage. Nodes will be added to the cluster transparently, with a simple boot, which will make the node amenable to a migrated process. Since the nodes are thin, version skew will be eliminated, as will traditional cluster installation (putting an image on every node).

The HPC cluster of the future should be “open from the ground up,” because it will be based on LinuxBIOS technology. The ramifications of putting Beowulf 2 technology such as Scyld on top of a LinuxBIOS platform will be very close to the HPC appliance.



IBM services for Linux

This chapter introduces you to the IBM services that help Linux customers gain the greatest advantage of a Linux-powered IBM solution. This chapter is intended to help those who are interested in Linux from technical specialists to sales, brand, marketing, and product area professionals. It provides the first fully consolidated reference to the continuously growing IBM information base on Linux.

The resources portrayed here provide a good starting point to more detailed information. They will help you to stay current with Linux's enhancements and new services, which continue to grow.

Some of the information presented in this chapter may change due to normal product life cycle variations. This is also true for products other than Linux. This means that even though a product may appear to be supported in this chapter, IBM may not support a specific product. Therefore, we encourage you to investigate such topics as end-of-service dates, product withdrawals, and other support restrictions while gathering information from this chapter.

Important: The services portrayed here are for IBM customers only, and may not be available in some specific countries or regions. Contact your local IBM Global Services representative for details.

11.1 IBM Global Services for Linux

IBM Global Services is the largest services organization with more than 140,000 employees in 164 countries or regions. It has skills professionals to help the customer unlock the business value of their information technology (IT) investments and become on demand businesses.

When IBM made the commitment to Linux and open source several years ago, it leveraged the expertise of thousands of strategic consultants, IT architects, and technical teams. IBM Global Services professionals combine years of experience on complex customer engagements in all environments with unique Linux-based skills. In doing so, they can help customers to assess, plan, design, implement, run, and support their Linux infrastructures, and solve business problems with advanced technology.

To learn more about IBM Global Services, see:

<http://www.ibm.com/services/>

You can choose to work with IBM Global Services to complete solutions. Or, you can select the service you need from among the industry's most comprehensive sets of Linux services.

11.1.1 Assess, plan, and design

IBM Business Consulting Services

With the acquisition of PricewaterhouseCoopers Consulting and its integration into IBM Business Consulting Services, IBM created an organization with unmatched industry insight, strategy process, deployment, and start-to-finish accountability. From business transformation to enterprise solutions in Supply Chain Management (SCM), Customer Relationship Management (CRM), and human capital, IBM helps you change “business as usual” to “on demand business.”

For more information about IBM Business Consulting Services, see:

<http://www.ibm.com/services/bcs/>

Linux strategy workshop

For leading e-businesses that need to speed integration and implement new applications across a multi-vendor system, Linux is worth investigating. It offers flexibility, dependability, and value. IBM Global Services has developed a customized engagement that help you initiate or expand Linux in an environment. At the end, you will walk away with a high-level strategy for

incorporating Linux that is based on your IT environment, your business issues, and your budget requirements.

For more information, see:

<http://www.ibm.com/services/its/us/linuxworkshop.html>

Server consolidation survey and other TCO and ROI assessments

IBM Global Services professionals use proven methodologies and tools to help customer determine the best solution for their infrastructure issues. For example, the server consolidation study is designed to provide a recommendation to optimize your infrastructure to bring an effective business solution for your e-business, enterprise resource planning (ERP), SCM, and business intelligence.

For more information, see:

<http://www-1.ibm.com/services/its/us/techassess.html>

Application innovation services

These services offer consulting and delivery expertise in end-to-end systems integration and custom application development in Linux environments. They are designed to drive innovation and expansion into new markets while reducing overall costs.

For more information, see:

<http://www.ibm.com/services/bis/integration.html>

11.1.2 Implement and run

Linux application services

There are two Linux application services of which you can take advantage.

Application porting

A range of services is available to migrate applications to the new Linux environment. This includes critical assessments, porting to Linux, application development when required, and application management and integration.

Application management services

These services encompass the full application life cycle: development, management, maintenance, support and governance for customer, legacy package, and e-business systems. Delivered onsite or through a global network

of IBM centers, they can address single, multiple, or an entire portfolio of applications.

For more information, see:

<http://www.ibm.com/services/ams/amsoffering.html>

System migration

IBM Global Services offers a comprehensive set of systems migration and implementation services for a heterogeneous environment. This section lists a sample of those that are available.

Migration services

Whether you have an IBM or other vendor environment, our migration specialists combine proven methodologies with strong project management in a team approach to help you minimize costly downtime or data loss. From assessment through deployment, we look at the total picture to help ensure your migration is thorough, timely, and successful.

For more information, see:

<http://www.ibm.com/services/its/us/migration.html>

Linux performance services

These services range from consulting to testing, and capacity planning to help you maintain an efficient and effective Linux infrastructure.

For more information, see:

<http://www.ibm.com/services/its/us/perftest.html>

Cluster implementation services

For pSeries and xSeries clusters, IBM Global Services can help you plan, install, support, and configure your solution.

For more information, see:

<http://www.ibm.com/services/its/us/drmk1m09.html>

File and print workload consolidation on zSeries

This service helps you consolidate a file, print, and Web-serving workload on Linux, to reduce costs and optimize operating efficiencies.

For more information, see:

<http://www.ibm.com/services/its/us/drmk1m01.html>

Site and rollout services

These services help you to make effective changes to your physical IT environment with reduced impact on your operations. When you move to a Linux environment, we can help you plan, design, and implement necessary infrastructure changes. This may include data center ventilation, cable and wiring layout, equipment modification, and more.

For more information, see:

<http://www.ibm.com/services/its/us/siteserv.html>

IBM DB2, WebSphere, implementation services

Our professionals can help you plan, install, and configure these key software components in your Linux environment.

For more information, see:

<http://www.ibm.com/services/its/us/swsupport.html>

Linux security services

The collection of Linux security services includes:

- ▶ **Enterprise security architecture:** Provides a review of your organization's overall business strategy, security policies, and environment to provide a complete picture
- ▶ **System security assessment:** Helps to identify vulnerabilities that may exist on your Linux operating system, platforms, and core middleware components. Both technology and management controls are reviewed.
- ▶ **System security implementation:** Services are designed to harden an existing security environment.
- ▶ **Ethical hacking:** This service provides after-the-fact vulnerability testing.

Linux technical training services

IBM Learning Service provides courses for Linux on IBM DB2, system administration, iSeries, zSeries, certification and Red Hat Linux, as well as Redbooks and free Web tutorials.

For a complete list of courses, see:

<http://www.ibm.com/services/learning/us/catalog/linux/>

IBM managed hosting: Linux virtual services

These services provide a Linux-based dynamic, on demand hosting environment. They allow the customer to pay, on a monthly basis, for as much

server, storage, and networking capacity as they need. This simplifies their capacity planning, while reducing their time to market, and avoiding up front capital expenses.

11.1.3 Support: IBM Support Line

Linux support is consistently one of the most cited requirements for customers. IBM offers “enterprise-grade” support for Linux through the Internet and voice support 24 hours-a-day, 7 days-a-week. This support ranges from answering usage questions to identifying problems, all revolving around the customer needs.

Support Line can either operate as a standard offering, or be customized with several options. With the strong alliance agreements IBM Global Services has in place with Red Hat, SuSE, and UnitedLinux, you have access to skilled help for supported Linux distributions.

IBM Operational Support Services: Support Line for Linux

This fee service provides remote Linux technical assistance to answer usage, installation, and technical questions. It also addresses product compatibility and interoperability concerns. And, it isolates and diagnoses problems. IBM Linux engineers are available 24 hours-a-day to provide defect support (fixes).

Support Line provides consistent, cross-platform Linux support for IBM @server platforms and Intel/AMD original equipment manufacturer (OEM) hardware. With this service, customers receive telephone access and electronic access to IBM services specialists.

Support Line is a single, integrated, remote support solution for Linux that provides:

- ▶ An 800 toll-free number, voice and electronic support
- ▶ Support for SuSE LINUX, Red Hat Linux, and UnitedLinux
- ▶ 8 a.m. to 5 p.m. support Monday™ through Friday in your time zone
- ▶ Full shift (24-x-7 optional coverage)
- ▶ Response time: Two hours prime shift, two to four hours off shift
- ▶ Available on zSeries, iSeries, pSeries, and xSeries servers, plus Intel/AMD OEM hardware
- ▶ Defect support by the IBM Change team to the source code level
- ▶ Interpretation of product documentation
- ▶ Configuration samples of IBM and multi-vendor database searches
- ▶ Planning information for software fixes

IBM Support Line for Linux offers:

- ▶ 24x7 enterprise-level remote support

- ▶ Fast and accurate problem resolution
- ▶ Access to IBM skilled services specialists
- ▶ Answers the “how-to” questions
- ▶ Resolutions to complex defect issues
- ▶ Support across four major distributions of Linux
- ▶ Remote assistance available through electronic and telephone access
- ▶ Electronic support and problem submission that optimizes your time and gives full online tracking of open support items
- ▶ Support Line options

IBM Operational Support Services: Support Line for Linux Clusters

This fee service provides remote Linux cluster technical assistance to answer usage, installation, and technical questions. It addresses product compatibility and interoperability concerns. And it isolates and diagnoses Linux cluster problems. It provides specialized Linux cluster support with open source and OEM. It offers remote problem identification and technical support for your Linux clustered environments. This service complements your in-house skills by providing comprehensive knowledge of Linux cluster environments.

This fee service provides:

- ▶ How-to, usage, defect, fix information
- ▶ An 800-toll free number and electronic access
- ▶ 8x5 or 24x7 coverage
- ▶ Defect support by the IBM Change team to the source code level
- ▶ Response time: Two hours prime shift, two to four hours off shift
- ▶ Available on Cluster 1300 and IA/32

Integrated cluster support is available for:

- ▶ IBM hardware
- ▶ Compute nodes
- ▶ System management and storage nodes
- ▶ OEM hardware bundled with cluster:
 - Myrinet and Cisco cluster/network switches
 - KVM switches and terminal servers
 - Power management modules
- ▶ IBM, OEM, and open source software including:
 - Cluster Systems Management (CSM)
 - General Parallel File System (GPFS)
 - Message Passing Interface (MPI) modules
 - Linux, and other open source software as supplied

IBM Operational Support Services: Advanced Support

IBM's Advanced Support provides the highest level of remote support available through IBM. The customer receives 24x7 coverage, 365 days-a-year. This service is tailored to meet unique customer needs for continuous, business-critical Linux system operation.

IBM operational support services: Custom Linux support

IBM offers remote Linux engineering support for creating modified versions of Linux. This services provides full remote support once you implement this custom distribution throughout your enterprise or customer base. By leveraging IBM's vast industry of Linux knowledge, IBM helps you to avoid serious pitfalls that can occur when you create your own distribution. We help you understand the necessary interdependent packages and required modules.

This service offers:

- ▶ Technical consulting: Development-level Linux support for Linux kernels and operating systems
- ▶ Tailored solutions to meet your development needs and schedules
- ▶ Development-level consulting in determining the appropriate Linux components, open source code, Linux tools, and applications for custom builds
- ▶ Compatibility and interoperability development support
- ▶ Assistance with customize Linux images or builds
- ▶ Assistance with custom Linux integration build
- ▶ Assistance with custom integration build testing
- ▶ 8 a.m. to 5 p.m. Monday through Friday availability
- ▶ Dedicated project team for specific skill area
- ▶ Collaboration with IBM development teams, Linux Technology Center (LTC), and the open source community
- ▶ Help to maximize mission-critical Linux system availability
- ▶ A focus on diagnostics, remote diagnostics, management, monitoring, and security
- ▶ Proactive development-level technical support based on an understanding of your mission-critical environment

IBM has the critical mass to deliver support teams in multiple worldwide locations. It has used some of its best talent in multiple locations to create the Change team. IBM can draw on the skills of over 200 key members of the Linux Technology Center found in more than 20 locations worldwide.

You can learn more about the IBM Support Line at:

<http://www.ibm.com/services/its/us/supportline.html>

11.2 IBM Products and Solutions Support Center

The IBM Europe Middle East Africa (EMEA) Products and Solutions Support Center (PSSC) is the place to convince customers, partners, and influential leaders about the advantages of IBM servers, storage platforms, and software in the support of their business critical solutions. As part of the Advanced Technology Solutions (ATS) organization, the PSSC ensures that your situation is always best served by the team of 170 highly skilled professionals with potentially external experts. Over 5,000 visitors, representing 800 companies from 55 countries, participated in the 950 engagements performed at the PSSC in 2002.

The main Web site for the PSSC is:

http://www.ibm.com/partnerworld/pwhome.nsf/pi/tech_support_pssc.html

The objective of the PSSC is to win business by showing the advantages of IBM @server and IBM TotalStorage solutions in support of customer or partner critical requirements. Where performance, scalability, availability, manageability and functionality are the issue, the 170 skilled professionals of the center give advice in defining the suitable architecture and proof scenarios. They also manage and implement the benchmark project to deliver convincing figures in line with customer expectations.

The support scope is not limited to system infrastructure. The years of experience that the PSSC ATS teams have built result in unique skills that are applied to the configuration and performance of e-business; business intelligence; EAS, CRM, and SCM; and scientific and technical applications on IBM server and storage platforms. This software and solution scope also encompasses Linux, the IBM application framework for e-business, and key ISV solutions in the EAS or e-business segments.

The core activity focuses on benchmarking, performance and sizing validation, and the New Workload Assistance program. This program offers zSeries expertise to leading-edge customers who sign up for early deployment of technologies in the areas of e-business, applications, and Linux. The PSSC infrastructure is also open to the early steps of the sales cycle for demand-generation events, tailored technical briefings, and workshops to convince your customer or partner of the value of the IBM proposition.

In March 2003, the PSSC was enhanced with the Storage Solutions Benchmark Center in addition to the existing center activities:

- ▶ Benchmark and training
 - pSeries (RS/6000) benchmarking experience since 1995
 - zSeries (S/390) since the opening of the center 10 years ago
 - xSeries (Netfinity) support since 1998
 - iSeries support since 2000
 - Testing Services Center 2000
 - Storage Solutions Benchmark Center in March 2003
- ▶ Pre-sales support
 - zSeries Solutions and New Technology Center (NTC)
 - EMEA Design Center for e-business on demand for all platforms in June 2000
 - EMEA Oracle IBM Joint Solutions Center (JSC) in June 1998 (xSeries, pSeries and zSeries)
 - Linux Competency Center in June 2001
 - EMEA Siebel IBM Support Center (ESISC) in April 2002
 - Grid Innovation Center in April 2002
- ▶ Services delivery
 - Education
 - Porting Center
 - GDPS® Solution Test in 2002
 - IBM eServer Executive Briefing Center (Technical and Marketing Events)
 - Business Recovery Services (BRS) and platform loan

Server computing power

As an average capacity:

- ▶ pSeries, iSeries, and xSeries servers:
 - More than 400 pSeries and SP nodes, 15 Teraflops
 - All IBM-supported SP and symmetric multiprocessor (SMP) hardware including the industry leader pSeries Model 690
 - Fifty-five TB of external and ESS storage
 - iSeries 10 servers, 110,000 CPW=1,000,000 TPC/C
 - xSeries 100 servers, high-end xSeries 440, xSeries 335 clusters, blades including EXA & FASTT
 - 300 PC NetVista

- ▶ zSeries
 - zSeries 800 and 900 latest CMOS technology with more than 82,000 Mips, 60 TB
 - Parallel sysplex architecture (coupling facility, sysplex timers)
 - Storage units: ESS, latest technology and more than 60 TB
 - Network units: 3745/3746, 3172, IBM Cisco routers, SAN Gateway, network attached storage (NAS) switch and more.
 - All our servers are part of a fully integrated production infrastructure allowing us to build heterogeneous infrastructures
 - Heterogeneous network security lab
 - IBM TotalStorage, SAN Gateway, and Storage Area Network (SAN) switch
 - NAS architecture
 - Heterogeneous server architecture and solutions

The total PSSC server capability heavily depends on running activities, covering four areas:

- ▶ zSeries: 82,000 Mips 60 TB
- ▶ pSeries: 15 Teraflops 55 TB
- ▶ iSeries: 10 high-end servers
- ▶ xSeries: 100 servers

All IBM servers are usually fully configured in terms of processors, memory, and I/O.

Skills

The available skills for architecture, installation, customization, maintenance, and performance among the 170 PSSC professionals include:

- ▶ Applications

Baan, Business Intelligence, Business Object, Citrix Metaframe, DataMining, e-business, Lotus Domino, WebSphere, ERP/SCM/CRM products, grid architecture (Platform, Globus, Avaki, GridSystems), Intelligent Miner, JD Edwards, Life Keeper (Linux Clustering), Microstrategy, MSCS (Windows Clustering), Oracle e-Business Suite (11i), PeopleSoft, SAP, Siebel, Windows 2000 (Server, Advanced Server, Data Center), Windows Server 2003, Windows Terminal Server
- ▶ Operating systems and subsystems

AIX, CICS, CSM, IBM Director, IMS, Linux, Linux for zSeries, Novell Netware, OS/400, PSSP, Server Management Software, UNIX, VMware ESX Server,

VM/ESA, Windows 2000 TSE and Citrix Metaframe, Windows 2000 and 2003, Windows XP, z/OS, z/OS UNIX System Services, Linux on zSeries, z/VM

► Architecture

Autonomic computing, availability, GDPS, grid computing, HACMP, Linux Clustering, Oracle Real Application Cluster (RAC), performance tuning, PPRC, pSeries/xSeries clustering, scalability, server consolidation and pSeries, Sysplex, IBM TotalStorage architecture, IBM Tivoli Storage Manager, XRC

► Network

ATM, FDDI, Advanced Networking (Load Balancing, Content Switching), Fiber Optics, GIGA Ethernet, Myrinet 2000, Network Management (Netview, Tivoli), Optical Networking using DWDM, (Dense Wave® Division Multiplexing) for metropolitan area network (MAN), Quadrics, SAN, SNA, TCP/IP, T/R, wireless networking, stress and load products (Mercury, Rational, Compuware)

► Relational database management system (RDBMS)

DB2 Universal Database (UDB) for z/OS and OS/390, DB2 UDB, Exchange, Informix, Lotus Domino, Microsoft SQL Server, Oracle/Oracle9i RAC

► Scientific and technical

Code porting and tuning, Fortran and C programming, parallelization on distributed and shared memory systems

Customer facilities

Despite its impressive computation power and its 6 000 m² (square meters) surface, the PSSC visitors never feel that they are “just another” partner to IBM. Instead, they have the impression that the entire infrastructure is tailored to their needs. This is due to the special care that the IBM team dedicates to adapting our offerings to each individual request and qualifying the expectations of visitors. This is also a results of the modularity of our welcome infrastructure with top class reception rooms, fully equipped for multimedia presentations, and modular technical activities rooms.

The details of the customer facilities are outlined here:

► 2000 m² customer reception area

- One 145 seats auditorium
- 6 briefing rooms of various sizes for 15 to 45 participants
- Rest and meeting area

► 2500 m² exhibition and server rooms and showroom

- One server exhibition lobby
- Two server and solution showrooms
- Two server and storage showroom (mid-range and high-end)
- Four server high computation rooms (system rooms)
- ▶ 1500 m² technical working rooms
 - IBM TotalStorage interoperability lab
 - Twenty-five benchmarking rooms of various sizes for teams from 3 to 30 participants
 - 10 education and workshop hands-on rooms

Contact information

EMEA IBM @server Executive Briefing Center

For further information about IBM @server Executive Briefing Centers, refer to the following Web site:

<http://www.ibm.com/servers/eserver/briefingcenter/>

This site provides links to all the IBM Executive Briefing Centers around the world including Austin, Böblingen, Greenock, Mainz, Montpelier, Poughkeepsie, Research Triangle Park, Rochester, San Jose, and Tucson. The Executive Briefing Center in Montpelier, France, is also the location for various benchmarking centers, solution centers, technology centers, and education centers. Some contact information is provided in the following paragraphs.

These centers and the personnel who staff them are in a state of continual evolution. Therefore, the following contact names may change over time. Your local IBM representative should be your first point of contact for queries about obtaining the services of these centers.

EMEA iSeries, pSeries, xSeries, zSeries Benchmarking Centers

For further information about the EMEA and other pSeries Benchmarking Centers, see the following Web page, and select the appropriate country:

<http://www.ibm.com/servers/eserver/benchmark.html>

In general, all benchmarking centers require you to contact your local IBM representative. This person makes arrangements with the benchmarking center for your benchmark project, including the fee schedule for your project.

pSeries and zSeries education

For further information about pSeries and zSeries education, contact:

Thierry Guffroy
 Telephone: 33 (0) 467 34 6871
 e-mail: mailto:thierry_guffroy@fr.ibm.com

Storage Solutions Benchmark Center

For further information about the EMEA Storage Solutions Benchmark Center, contact the Executive Briefing Center at Montpellier, France.

Eric Dambroise
Telephone: 33 (0) 467 34 4988
e-mail: <mailto:dambrois@fr.ibm.com>

Scientific & Technical Computing Services

For further information about Scientific & Technical Computing Services, contact:

Jean-Yves Leclere
Telephone: 33 (0) 467 34 4619
e-mail: <mailto:leclere@fr.ibm.com>

zSeries & Solutions New Technology Center

For information about zSeries & Solutions New Technology Center, contact:

Patrick Joussen
Telephone: 33 (0) 467 34 6605
e-mail: <mailto:joussen@fr.ibm.com>

EMEA Design Center for e-business on demand

For information about EMEA Design Center for e-business on demand, contact:

Alain Poquillon
Telephone: 33 (0) 467 34 6173
e-mail: <mailto:apoquillon@fr.ibm.com>

Grid Innovation Center

For further information about the Grid Innovation Center, contact:

Yann Guérin
Grid TSM
Telephone: 33 (0) 467 34 4252
e-mail: mailto:yann_guerin@fr.ibm.com

EMEA Siebel IBM Support Center (ESISC)

For information about EMEA Siebel IBM Support Center (ESISC), contact:

François-Rene Rougeaux
Telephone: 33 (0) 467 34 6754
e-mail: <mailto:rougeaux@fr.ibm.com>

IBM Oracle Joint Solutions Center (JSC)

For further information about IBM Oracle Joint Solutions Center (JSC), contact:

Alain Benhaim
Telephone: 33 (0) 467 34 4961
e-mail: <mailto:benhaim@fr.ibm.com>

Porting Center

For information about Porting Center, contact:

Eric Dambroise
Telephone: 33 (0) 467 34 4988
e-mail: <mailto:dambrois@fr.ibm.com>

GDPS Solution Test

For further information about GDPS Solution Test, contact:

Daniel Roman
Telephone: 33 (0) 467 34 6034
e-mail: mailto:Daniel_Roman@fr.ibm.com

Operations & Infrastructure Center Manager

For information regarding IBM @server & Network Infrastructure, BRS, loan, and try and buy activities, contact:

Antoine Guillamo
Telephone: 33 (0) 467 34 6558
e-mail: mailto:antoine_guillamo@fr.ibm.com

Any Request

For any other information about the EMEA PSSC, call or send e-mail to the contact listed here:

Telephone: 33 (0) 467 34 4509
e-mail: <mailto:justcall@fr.ibm.com>

For information about other IBM centers, call your local IBM representative.

11.3 IBM Software Services for Linux

IBM is rolling out services to help customers rapidly deploy and optimize Linux-based e-business environments. IBM services for Linux covers Software Group products such as WebSphere, Data Management (DB2), Lotus, and Tivoli. IBM is also delivering migration services to help customers port applications to a wide variety of IBM and non-IBM platforms running Linux.

11.3.1 WebSphere

Many customers need WebSphere as a front end to their legacy applications. Customers need to access legacy applications (for example, CICS) through a browser interface to provide users (internal or external) with Web-enabled access to multiple existing applications. They also want solutions that do not require a long development cycle and that support an open-standards compliant integrated infrastructure. Customers want to leverage existing transactional applications, database assets, and existing investment in hardware platforms that have a superior scalability characteristic on which to run their new application.

Every customer with an IBM mainframe capable of having Integrated Feature for Linux (IFL) presents a good fit for WebSphere for Linux. IFL-capable systems are the 9672 G Series (G5 and G6), Multiprise 3000, zSeries 800, and zSeries 900 vendors.

IBM provides a variety of services for WebSphere that go from migration to specific training in all WebSphere products. For WebSphere software services and training, see:

<http://www.ibm.com/software/os/linux/software/websp.shtml>

For more information about WebSphere training and technical enablement, see:

<http://www.ibm.com/developerworks/websphere>

This Web page takes you to the developerWorks portal into WebSphere. Here you will find training and technical enablement as well as links to other WebSphere services, products, and reference information.

IBM also provides support for the wide range of WebSphere products made to fit the Linux environment. Table 11-1 lists references where you can find support for WebSphere products that run on Linux. As with all lists of products and Web sites, the information presented in this table may change and evolve over time. If you are unable to find a particular link, try to back up to a primary link such as the links for `software.ibm.com/websphere` or `developerworks/websphere`.

Table 11-1 WebSphere for Linux product services matrix

WebSphere services	Service solutions	More informations
WebSphere Application Server	Installation Services for WebSphere Application Server	http://www.ibm.com/services/its/us/ss-websphere.html
WebSphere Edge Server	Implementation Service for WebSphere Edge Server	http://www.ibm.com/services/its/us/ss-websphredgesvr.html

WebSphere services	Service solutions	More informations
WebSphere Application Server	Migration Services for WebSphere Application Server	http://www.ibm.com/services/its/us/migrationwebappserv.html
WebSphere Portal Server	Installation Services for WebSphere Portal Server	http://www.ibm.com/services/its/us/implementationwebmulti.html
WebSphere MQ	Business Integration Services	http://www.ibm.com/software/integration/websphere/services/
WebSphere Application Server	WebSphere on Demand Application development skill transfer	http://www7.software.ibm.com/vadd-bin/ftpd1?1/vadc/wsdd/pdf/services/042903_WASOnDemandADSkillsTransfer.pdf
Web services	IBM Services for Web Services	http://www7.software.ibm.com/vadd-bin/ftpd1?1/vadc/wsdd/pdf/services/031803_ServicesforWebServices.pdf

11.3.2 DB2 for Linux

DB2 for Linux Version 8.1 is now available. You can learn more about this release on the IBM DB2 for Linux Web site at:

<http://www.ibm.com/software/data/db2/linux/validate/>

You should also refer to Chapter 9, “IBM Software for Linux” on page 265, for DB2 for Linux hardware and software requirements.

IBM offers the following services related to DB2 for Linux:

- ▶ **Remote DB2 Installation:** This service is supported for DB2 for Linux running under a logical partition (LPAR), a guest system (VM), or a stand-alone Linux. For more information, see:
<http://www.ibm.com/services/its/us/drmk1m12.html>
 IBM can also help customers to migrate to DB2 for Linux. Learn about a variety of options and helpful information on the Web at:
<http://www.ibm.com/software/data/db2/migration/>
- ▶ **DB2 for Linux support:** For technical support, product information, and all other DB2 for Linux support-related issues, go to:
<http://www.ibm.com/software/data/db2/udb/support.html>
 For information about DB2 consulting services and other services, refer to:
<http://www.ibm.com/services/>

11.3.3 Tivoli

Customers' growing use of Linux systems within their Tivoli-managed environments has extended the commitment of IBM to scale Linux services and support to Tivoli products.

IBM announced Tivoli Linux enablement since the summer of 2002 for its security software and Web management products. For more information, see:

<http://www.ibm.com/software/tivoli/solutions/linux/>

Linux enablement was also extended to Tivoli Storage Management, Configuration, Operations, Performance, and Availability products. For details about the current Linux supported Tivoli products, refer to:

<http://www.ibm.com/software/tivoli/>

You can access services for Linux supported Tivoli products, which include consulting, training, and certification on the Tivoli services Web site at:

<http://www.ibm.com/software/tivoli/services/>

Tivoli's Customer Support is quite extensive and complete. Some services are available only for registered users. You can find Tivoli Customer Support on the Web at:

<http://www.ibm.com/software/sysmgmt/products/support/>

Product support matrices for Tivoli

There is a wide array of Tivoli products for all platforms. This Web site is a portal to Tivoli products for Linux:

<http://www-306.ibm.com/software/os/linux/software/tivoli.shtml>

In addition, there is a Tivoli product cross-reference matrix at:

<http://www-306.ibm.com/software/tivoli/products/product-matrix.html>

This matrix is a reference to the changes in Tivoli product names. It may be a useful reference for sorting out the array of Tivoli products.

11.3.4 Lotus

To support and enhance your Notes and Domino environment, Lotus offers a full range of professional services including consulting, education, and customer support. See the IBM Lotus Software site for more information:

<http://www.lotus.com/services/education.nsf/wdocs/serviceshomepage>

11.4 IBM platform services for Linux

The xSeries, zSeries, iSeries, and pSeries platforms work with Linux as an operating system today, among the other operating systems that exist in the marketplace. As software-based services for Linux are needed, so is a well-founded, platform-based service and support structure that is still growing.

The IBM Linux Support Line provides operational support services and a premier remote technical support service. For more information, see:

<http://www-1.ibm.com/services/its/us/supportline.html>

IBM provides technical support for the major distributions of the Linux operating system running in the xSeries, zSeries, iSeries, and pSeries platforms. IBM also provides support for some non-IBM applications that operate in a Linux environment. IBM helps answer how-to questions, performs problem source determination, and provides mechanisms for a solution. In addition, by leveraging partnerships with the key distributors of the Linux operating system, IBM provides defect-level support for the Linux operating system. Remote assistance is available through toll-free telephone or electronic access depending on the country or region.

IBM provides services for all currently supported xSeries, zSeries, iSeries, and pSeries platforms in a different degree of variety and complexity through its Linux Portal. For more information, see:

<http://www.ibm.com/linux/>

IBM also provides Learning Services for customers who are interested in knowing more about Linux and making it their key operating environment for e-business. IBM Linux education and training services offer a full curriculum and help to establish a personal Linux education plan. For details about IBM Linux Learning Services and Linux training, you can visit:

<http://www.ibm.com/services/learning/us/catalog/linux/>

11.4.1 zSeries services

Table 11-2 lists various resources for platform-based information about services that are available for the zSeries server.

Table 11-2 zSeries Linux information matrix

Information	Location on the Web
Linux for zSeries applications	http://www.ibm.com/servers/eserver/zseries/solutions/s390da/linuxisv.html
Linux for zSeries operating system download	http://www.ibm.com/servers/eserver/zseries/os/linux/dist.html
Linux for zSeries development tools	http://www.ibm.com/servers/eserver/zseries/os/linux/ldt/
Linux for zSeries library	http://www.ibm.com/servers/eserver/zseries/os/linux/press.html
Linux for zSeries education	http://www.ibm.com/servers/eserver/zseries/os/linux/ed.html

11.4.2 iSeries services

IBM provides a variety of services for customers interested in consolidating their servers on the iSeries server. To learn more about these services, see:

<http://www.ibm.com/servers/eserver/series/linux/adv.html>

Table 11-3 lists various references for platform-based information about services that are available for the iSeries server.

Table 11-3 iSeries Linux information matrix

Information	Location on the Web
Linux for iSeries applications	http://www.ibm.com/servers/eserver/series/linux/apps.html
Linux for iSeries education	http://www.ibm.com/servers/eserver/series/linux/educ.html
Linux for iSeries planning	http://www.ibm.com/servers/eserver/series/linux/plan.html
Linux for iSeries developer resources	http://www.ibm.com/servers/eserver/series/linux/resc.html

Information	Location on the Web
Linux for iSeries reference library	http://www.ibm.com/servers/eserver/series/linux/library.html
Linux for the iSeries support	http://www.ibm.com/servers/eserver/series/linux/support.html

11.4.3 pSeries services

The pSeries server makes a reliable foundation for fast 64-bit native Linux servers. See Chapter 6, “pSeries and Linux” on page 173.

IBM provides hardware and software support service for the new pSeries 630 6C4 and 630 6E4 Linux-ready express configurations that are now available. IBM will provide more services for Linux for pSeries in response to growing interest from customers to use Linux as their native operating system in this platform.

For more information, see the IBM Linux for pSeries Web site at:

<http://www.ibm.com/servers/eserver/pseries/linux/index.html>

11.4.4 xSeries services

IBM is leveraging enterprise scalability and availability at attractive prices through xSeries Linux enablement. The combination of Linux with the xSeries by IBM has proven successful. Many resources and information are available to users who are interested in using Linux on the xSeries as their key environment to meet their needs.

Table 11-4 lists various references that provide platform-based information about services that are available for the xSeries server.

Table 11-4 xSeries Linux information matrix

Information	Location on the Web
Linux for xSeries	http://www.pc.ibm.com/us/eserver/xseries/index.html
IBM @server BladeCenter	http://www-1.ibm.com/servers/eserver/bladecenter/
Linux clustering on xSeries	http://www.ibm.com/servers/eserver/clusters/
Linux for xSeries ISV resources	http://www.ibm.com/developerworks/offers/linux-speed-start/isv.html

11.4.5 Consolidated platform matrix

For more valuable information about IBM Linux for all IBM @server services, as well as latest information, updates, new services available, etc., refer to the IBM Linux dedicated Web site for each platform as listed in Table 11-5.

Table 11-5 IBM Web site matrix for each Linux platform

Platform	Location on the Web
zSeries	http://www.ibm.com/servers/eserver/zseries/os/linux/index.html
iSeries	http://www.ibm.com/servers/eserver/iseries/linux/index.html
xSeries	http://www.ibm.com/servers/eserver/xseries/linux/index.html
pSeries	http://www.ibm.com/servers/eserver/pseries/linux/index.html

11.5 Linux assessment services

This is an example service offered by IBM.

11.5.1 Business challenge

Are you just introducing Linux? Are you already piloting Linux in your environment or expanding the use of open source? You need to know if your IT infrastructure and your business objectives are aligned to ensure that you have the right solution for your needs.

11.5.2 Solution

The skilled consultants and IT architects on the IBM Linux Assessment Services team can help you assess your environment and business goals. They use proven tools and methodologies, and leverage the expertise of hundreds of Linux engagements across all industries.

11.5.3 At a glance: Solution components

IBM Linux Assessment Services has a range of offerings and capabilities to assist you:

- ▶ Linux strategy workshop
- ▶ Grid innovation workshop
- ▶ Performance management and capacity planning assessments
- ▶ Return on investment and total cost of ownership (TCO) studies
- ▶ Server consolidation studies

- ▶ Security assessments: enterprise, system, application
- ▶ Migration and porting assessments

11.5.4 Benefits

With these services, you can gain a thorough understanding of the Linux strategy designed for your company, and the underlying architecture required for successful implementation.

11.5.5 Technical details and architecture

Assessments are consulting engagements in which the customer is an active partner and participant.

11.6 Server consolidation study

IBM Technology Assessment and Consulting Services server consolidation study provides a stable foundation for rapid deployment of current and future business solutions. This study is designed to bring a solution to your needs from hardware, software, data integration and application migration services, and more. It provides recommendations for optimizing your infrastructure to bring an effective business solution for your e-business, ERP, SCM, and business intelligence.

11.6.1 Highlights

This study provides a detailed picture of your current IT assets, system configurations, and utilization. The highlights of the study include:

- ▶ It assesses your current IT environment.
- ▶ It analyzes your IT resources to identify cost saving opportunities through consolidation.
- ▶ It recommends how to make your IT infrastructure more cost-effective and efficient.
- ▶ It furnishes a written report to assist in executive decision making regarding a future IT infrastructure.
- ▶ It provides experienced IBM IT consultants, specialists, and architects to help you succeed.
- ▶ It lowers TCO by optimizing existing IT infrastructure.
- ▶ It improves service levels by having continuous availability.

11.6.2 Pricing

This service is designed to be a custom priced offering. The IBM Technology Assessment and Consulting Services server consolidation study is custom-priced based on the number of servers in consideration for consolidation and the size of the overall project.

11.6.3 IBM Linux Implementation Services

This is a service offering available from IBM.

Business challenge

Installing new software or hardware systems can be time-consuming and complex, particularly when migrating to a new operating environment. You risk disrupting your business, draining internal resources, or spending time to acquire new skills.

Solution

IBM Linux Implementation Services solution provides a range of installation and implementation services to help customers get Linux up and running as quickly and painlessly as possible.

At a glance: Solution components

The solution components include:

- ▶ Installing IBM and IBM Hardware Storage
- ▶ Installing and configuring IBM software for Linux: DB2, WebSphere, Tivoli
- ▶ Installing custom Linux clusters, end-to-end
- ▶ Implementation of Linux security architecture
- ▶ Site enablement and rollout services
- ▶ Skill transfer to staff

Benefits

Customers can maintain availability for their customers, suppliers, and employees while reducing downtime and enhancing performance.

11.7 Examples of Linux services

IBM Global Services has been successful in delivering services to new IBM Linux clients. This section shares a few examples of such situations.

11.7.1 Example 1: Billion-dollar financial institution

This customer was running Red Hat 6.2 on a VA Linux hardware cluster solution. The customer's financial application would crash and core dump on 3% of the runs. There was no pattern to when the failure would occur. The problem plagued the customer's IT department for over two months. No other Linux service providers could identify the problem.

The same problem was reported sporadically on the kernel mailing list for over two years with no one able to identify the problem, nor offer a solution. Any attempts to instrument the kernel to isolate the problem changed the timing of the code and the problem would disappear (masked). The only possible way to find the root cause was through detailed code inspection that involved IA32 assembly code at their most fundamental level. Through a series of test cases and code inspection, IBM isolated the problem to a bug in `switch_mm()`, which resulted in a process running with a "stale" `ldt`.

IBM recommended a fix to the kernel mailing list, and it was accepted by Linus Torvalds. He responded in a personal note back to the IBM Linux Change team and released it in the v2.4.8 Linux kernel.

11.7.2 Example 2: Multi-million dollar advertising company

This customer was running Red Hat 7.1. They experienced a situation where an embedded Linux NFS client would error on its NFS mounts from the server at random points in time, but within a 24-hour period of booting. IBM reproduced the problem using the exact hardware and a copy of the software that the customer was running. Using iterative debug kernels and creating kernel dumps on error, IBM isolated the problem to NFS kernel code that errored with hard links beyond the exported directory.

IBM provided a work around immediately to the customer, while a patch was developed and consensus in the open source community was gained for a long-term code fix. Agreement was made and the fix will be included in the latest 2.4.x kernel. Previous threads to the kernel mailing list indicated that others had run into this problem previously, but no one could isolate it to the root cause.

11.7.3 Example 3: One of the largest network services providers

When this customer, running Red Hat 6.2, upgraded their hardware, they ran into a problem where the system hung 15 minutes after boot time. The original problem statement indicated that no software changes were made and that only processor upgrades were completed from one release to another. Working through the weekend to investigate the problem, IBM found that a boot

parameter was set to hard code memory boundaries. This was in conflict with reserved space for the BIOS.

Further investigation demonstrated that the boot parameter did not need to be set because the system properly probes memory size and sets it accordingly. In Red Hat 7.1, code was added to check BIOS memory allocation and not to override it with boot-time configuration settings.



Linux solutions

This chapter provides examples of solutions where IBM brings together a combination of IBM products and services, Business Partner relationships, ISV applications, Linux, and other open source components to solve a critical business challenge. They are all backed up with customer success stories, and all ready to work for our customers.

12.1 Total IBM Linux solutions abstracts

The Linux at IBM Solutions Web page lists many notable Linux-based solutions, utilizing IBM hardware. The majority of solution pages contain real-world case study appendages. In addition, each page contains many hotspots, taking the user to locations that contain further details about the solution or referred to companies.

The following examples are drawn from the IBM Linux Solutions Web site. For a complete listing of current solutions visit:

http://www.ibm.com/linux/va_4049.shtml

These selected solutions provide a good cross-reference. They illustrate what IBM Business Partners and Customers are accomplishing with Linux. To simplify this section, we include an abstract for each solution, so you can quickly understand the impact Linux makes in the marketplace. We also list the corresponding Web address back to the provided page.

12.1.1 IBM Global Solutions Directory

IBM offers a Global Solutions Directory (GSD), storing IBM @server proven solutions, implemented by IBM Business Partners. The GSD forms the IBM @server Solutions Connections site. This allows for easy access by IBM customers or business partners, and for each to make informed business and technology decisions.

Tip: The IBM @server Solutions Connection site offers a customer marketing sheet that illustrates the benefits provided through this service. For more information, see:

<http://www.developer.ibm.com/welcome/eserver/eSC/CSFServlet?mvcid=Welcome&packageid=000>

<http://www.ibm.com/servers/solutions/finder/graphics/common/pdf/custbrochure.pdf>

12.1.2 Computer-aided engineering

Business challenge

A company desires four improvements:

- ▶ Reduce the number of physical prototypes
- ▶ Decrease engineering cost

- ▶ Generate required data to meet safety standard
- ▶ Shorten time-to-market

Solution

IBM and our Business Partners have helped many customers successfully deploy computer-aided engineering solutions on Linux. The IBM Computer Aided Engineering solution runs simulation applications for crash test, computational fluid dynamics, along with noise, vibration, and harshness with higher performance. It does this at a cost less than those offered by traditional platforms. IBM and our Business Partners have helped many customers successfully deploy computer-aided engineering solutions on Linux.

At a glance: Solution components

The IBM Computer Aided Engineering Solution provides high performance computing to the automotive and aerospace industry. This solution is available worldwide to large enterprises and small to medium-size businesses. This solution includes the components listed in Table 12-1.

Table 12-1 Computer Aided Engineering Solution components

IBM product category	Included solution brand
IBM hardware	xSeries, IntelliStation®, and IBM Storage
Services and support	IBM Support Line, Technology Services
ISV applications	LS-DYNA from LSTC, CFD from EXA
Linux distributions (subject to change)	Red Hat 7.x and Enterprise AS 2.1

Benefits

From car manufacturers and suppliers to distributors and dealers, companies in the automotive industry are finding that their business models are challenged at every turn. Competition is fierce, global, and based on the efficiency of the entire supply chain. Rising costs can no longer be passed along to customers. Instead, pressures to reduce costs are being passed down the value chain, paring already thin margins. As automotive companies realize that they must be as innovative as ever, and more cost-efficient, they are re-examining traditional IT practices. For example, many companies are finding that it no longer makes sense to deploy and maintain core business applications on proprietary operating systems. In their search for an industry-standard, reliable, cost-effective platform, they are increasingly turning to Linux.

12.1.3 Linux in exploration and production

Business challenge

Increasingly volatile and unpredictable oil and gas prices, coupled with regional and global economic downturns, make the ability to find and commercially exploit oil and gas deposits even more critical. Information technologies have played a key role in exploration and production, important drivers of past success in the oil and gas industry. Now, with intensifying competition for prime exploration acreage and a greater need to increase recovery from existing fields, development and deployment of faster, more powerful technologies continues to be crucial.

Solution

IBM and Linux are providing cost-effective ways to address petroleum industry challenges. Leveraging the power of Linux on IBM hardware and leading edge ISV solutions help to build revenue and cut costs in many ways. This includes high-performance computing, streamlined operations, and distributed enterprise.

At a glance: Solution components

The Linux in exploration and production solution provides high performance computing to the petroleum industry. The Linux in exploration and production solution is available worldwide to large enterprises and small- to medium-size businesses. This solution includes the components listed in Table 12-2.

Table 12-2 Exploration and production solution components

IBM product category	Included solution brand
IBM hardware	xSeries, IntelliStation, Storage
IBM middleware	DB2, WebSphere, and Tivoli
Services and support	Support Line, Virtual services, Learning Services, workshops
ISV applications	Landmark Graphics
Linux distributions	Red Hat 7.x and Enterprise AS 2.1

Benefits

Linux improves performance and lowers the cost of high performance computing, thereby improving productivity. Lower cost and better flexibility of Linux in the exploration and production solution reduce the total cost of ownership (TCO) for petroleum customers.

On the technical side, the Linux Geocluster application, processes seismic survey data for oil companies benefitting businesses with 1.25 Teraflops of processing power at 50% of the cost of alternative solutions. The end result equates to more accurate images in *weeks* as opposed to months.

Technical details and architecture

The Linux in exploration and production solution uses Landmark Graphics running on a high performance Linux. Landmark Graphics provides seismic data processing, geophysical interpretation tools, and 4-D reservoir modeling.

Table 12-3 Exploration and production solution

Type	Location on the Web
Abstract	http://www-1.ibm.com/linux/solutions/eandplinuxsolution.shtml
Press release	http://www-3.ibm.com/software/success/cssdb.nsf/CS/NAVO-4YXQAW?OpenDocument&Site=linuxatibm
Brochure	http://www-1.ibm.com/linux/solutions/files/Linux_in_Petroleum_Brochure.pdf

Attention: The abstract contains a link to the *Linux in Petroleum Brochure.pdf*, which generates an error, due to an inaccurate hot spot link. Use the Brochure Web address in Table 12-3 to display the PDF file properly.

12.1.4 IBM Lotus Domino Web Access for Linux

Business challenge

Today’s businesses must run at a fast pace. The critical edge requires access to corporate messaging services, and personal information that helps your company communicate with your customers, so that your customers receive the information they need when they need it.

Solution

Help meet the demands of your users by providing critical messaging and personal information management (PIM) access when customers need it, regardless of where they are working. IBM Lotus Domino Web Access (iNotes) is a sophisticated Web client that gives users many of the messaging and collaboration features that were previously available only with a Lotus Notes client.

With this solution, you get the reliability and security features of an IBM Lotus Domino server delivered through a Web browser. Whether a user is at an

Internet kiosk or at a PC, whether online or disconnected from the network, that user can obtain ubiquitous, authenticated access to supported Lotus Domino data, including e-mail, calendar, to-do lists, personal contacts, and notebook.

Domino Web Access 6.5 extends messaging and collaboration to Linux clients with support for the Mozilla browser. This makes it the first-ever fully supported client-to-server collaboration solution for Linux.

At a glance: Solution components

This solution is available worldwide and across industries. It is ideal for small- to medium-sized businesses and enterprise. This solution includes the components listed in Table 12-4.

Table 12-4 Lotus Domino Web access solution components

IBM product category	Included solution brand
IBM middleware	Lotus Domino, Lotus Domino Web Access 6.5 e
Services and Support IBM Support Line	IBM Software Services for Lotus offers QuickStart in Asia-Pacific
Business partner services	Dependent on geographical location and industry needs
Linux distributions	Domino Web Access 6.5 supports Mozilla 1.3.1 on Red Hat Linux 8.0, 9.0, and SuSE 8.0 with the Domino server running on Red Hat Advanced Server 2.1 or United Linux 1.0 Sp2

Benefits

The benefits of this solution are that it:

- ▶ Delivers sophisticated Web messaging and PIM with an unmatched feature set and easy-to-use interface
- ▶ Provides offline support to answer e-mail or view calendar through replication
- ▶ Helps protect your environment and your information with comprehensive security features
- ▶ Helps reduce TCO through rapid, low-touch deployment, and with a minimal need for training
- ▶ Extends messaging and collaboration services to users who may not require a full Lotus Notes client
- ▶ Integrates with collaboration tools such as Lotus Instant Messaging and Lotus Team Workspace

Also, with the release of Lotus Domino Web Access 6.5, customers can benefit from a complete client and server Linux solution for enterprise messaging.

12.1.5 Real-time securities and risk analysis on Linux

Business challenge

There is a need for real-time access to information and analysis. This enables decision-making and trading with greater speed, efficiency, and risk management, all to better serve customers' high-volume transactions.

Solution

IBM Finance Foundation for Capital Markets is designed to offer a solid competitive advantage to financial institutions using specialized computing systems to analyze and manage financial risk, and cover positions on a real-time basis. This solution is an ideal platform for time-critical, compute-intensive financial transactions such as arbitrage and exotic trading. This solution is a combined in-memory and object relational solution that can handle up to 100,000 ticks per second on standard UNIX hardware. It is unlike offerings by proprietary in-memory-based storage, and analytics competitors and standard relational database competitors.

This solution is ideal for risk managers, quantitative analysts, brokers, hedge fund managers, market data managers, operations managers, and IT managers.

At a glance: Solution components

A worldwide solution, this solution includes:

- ▶ IBM Middleware, IBM Informix Dynamic Server, IBM Informix NAG DataBlade™ Module, IBM Informix TimeSeries DataBlade, and IBM Informix TimeSeries Real-Time Loader
- ▶ Services and Support IBM Linux Application Porting, IBM Linux Implementation Services, and IBM Support Line

This solution includes the components listed in Table 12-2.

Table 12-5 Real-time securities and risk analysis solution components

IBM product category	Included solution brand
IBM hardware	xSeries
IBM middleware	IBM Informix Dynamic Server IBM Informix NAG DataBlade Module IBM Informix TimeSeries DataBlade IBM Informix TimeSeries Real-Time Loader

IBM product category	Included solution brand
Services and support	IBM Linux Application Porting IBM Linux Implementation Services IBM Support Line

Benefits

The benefits are:

- ▶ **Speed:** When dealing with large datasets, this server-centric model ranges from 20 to 2,500 times faster than a relational database.
- ▶ **Volume:** This solution easily scales to loading rates above 100,000 ticks per second, while simultaneously running multiple queries.
- ▶ **Accuracy:** High-performance tools can cut through irrelevant detail and deliver precise information.
- ▶ **Efficiency:** Starting data in a time series may reduce a dataset to approximately half of its actual size and is stored contiguously on disk for fast access by the server.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 387.

- ▶ *IBM Tivoli Storage Manager Implementation Guide*, SG24-5416
- ▶ *IBM Tivoli Monitoring Version 5.1: Advanced Resource Monitoring*, SG24-5519
- ▶ *e-Business Intelligence: Leveraging DB2 for Linux on S/390*, SG24-5687
- ▶ *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765
- ▶ *Linux for WebSphere and DB2 Servers*, SG24-5850
- ▶ *Lotus Domino R5 for Linux on IBM Netfinity Servers*, SG24-5968
- ▶ *Linux Web Hosting with WebSphere, DB2, and Domino*, SG24-6007
- ▶ *Linux Applications on pSeries*, SG24-6033
- ▶ *CICS Transaction Gateway V5 The WebSphere Connector for CICS*, SG24-6133
- ▶ *A Practical Guide to Tivoli SANergy*, SG24-6146
- ▶ *IBM WebSphere V4.0 Advanced Edition Handbook*, SG24-6176
- ▶ *Linux System Administration and Backup Tools for IBM @server xSeries and Netfinity*, SG24-6228
- ▶ *e-Business Intelligence: Data Mart Solutions with DB2 for Linux on zSeries*, SG24-6294
- ▶ *IBM TotalStorage Enterprise Storage Server Model 800*, SG24-6424
- ▶ *Linux Application Development Using WebSphere Studio 5*, SG24-6431
- ▶ *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554
- ▶ *Linux Clustering with CSM and GPFS*, SG24-6601
- ▶ *Building a Linux HPC Cluster with xCAT*, SG24-6623
- ▶ *WebSphere Application Server Test Environment Guide*, SG24-6817

- ▶ *Lotus Domino 6 for Linux*, SG24-6835
- ▶ *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844
- ▶ *IBM WebSphere Portal V4.1 Handbook Volume 1*, SG24-6883
- ▶ *Up and Running with DB2 for Linux*, SG24-6899
- ▶ *IBM Lotus Domino 6.5 for Linux on zSeries Implementation*, SG24-7021
- ▶ *WebSphere Commerce V5.4 for Linux, Infrastructure and Deployment Patterns*, REDP0214
- ▶ *Linux on IBM zSeries and S/390: Securing Linux for zSeries with a Central z/OS LDAP Server (RACF)*, REDP-0221
- ▶ *IBM WebSphere Portal V4.1.2 in a Linux Environment*, REDP0310
- ▶ *WebSphere Application Server V4 for Linux, Implementation and Deployment Guide*, REDP0405
- ▶ *e-commerce Patterns for Linux on zSeries Using WebSphere Commerce Suite V5.1 Patterns for e-business series*, REDP0411
- ▶ *Deploying Apache on IBM @server BladeCenter*, REDP3588
- ▶ *IBM WebSphere V5.0 for Linux, Implementation and Deployment Guide - WebSphere Handbook Series*, REDP3601
- ▶ *Tivoli NetView Performance Monitor for TCP/IP 1.4 Using Linux as the Web Server*, REDP3602
- ▶ *Linux on IBM @server zSeries and S/390: z/VM Configuration for WebSphere Deployments*, REDP3661
- ▶ *Using IBM Tivoli System Automation for Linux*, REDP-3716

You can also find a complete listing of Linux-related IBM Redbooks on the Redbooks portal at:

<http://publib-b.boulder.ibm.com/redbooks.nsf/portals/Linux>

Other publications

These publications are also relevant as further information sources:

- ▶ *SWG Products, Solutions, and Strategies Reference Guide*, G325-2132
- ▶ *SWG Software Reference Guide*, Z325-5422 (IBM internal use only)

Referenced Web sites

These Web sites are also relevant as further information sources:

- ▶ IBM Linux Web site:
<http://www.ibm.com/linux/>
- ▶ The Linux Technology Center:
<http://www.ibm.com/linux/ltc>

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Numerics

10-Gigabit Ethernet 330
31-bit environment option 70
64-bit environment option 70

A

Advanced Power Management (APM) 57
AIX 5L 177
AIX and Linux 180
AIX Toolbox for Linux Applications 176
antivirus
 scanners 9
Apache 162
application solution 150
 iSeries 218
ARPAnet 49
AT&T 49
autonomic computing 348
availability 315

B

backup solutions 158
Bastille 34
BBDO Dusseldorf 345
Beowulf 334
Beowulf cluster 335
Beowulf cluster technology 192
Beowulf program 335
binaries 33
binary code 47
binary file 47
BIND (Berkeley Internet Name Daemon) 19
blade server 166
BladeCenter 166
BladeCenter HS20 166
Blue Gene research project 318
Blue Gene/L 318
BSD license 51
BSD UNIX 336
Business Intelligence (BI) 22

C

CES (Clustered Enterprise Servers) 324
channel bonding 335
client 23
 thick 23
 thin 23
closed software 49
cluster
 applications 332
 architecture 317
 autonomic computing 348
 availability 315
 benefits 314
 Beowulf 334
 bottlenecks 317
 Bproc 336
 Cluster 1350 324, 338
 Cluster 1600 323
 Clustered Enterprise Servers (CES) 324
 ClusterProven 342
 commodity hardware 314, 317
 components 324
 coupling 318
 distributed computing environment 318
 emerging roles 321
 failover 320
 Fast Ethernet 329
 FAStT200 330
 Fibre Channel 330
 file system 331
 General Parallel File System (GPFS) 332
 Gigabit Ethernet 329
 grid 346
 grid computing 346
 heartbeat 335
 high availability 142
 high availability cluster 335
 High Performance Computing (HPC) 316
 High Throughput Computing (HTC) 318, 332
 high-availability computing (HAC) 319
 high-performance computing 144
 history 334
 iSeries 212
 kernel extension 336

- Linux 142
- Linux high availability 335
- Linux Virtual Server (LVS) 336
- load balancing 143, 321, 333, 336
- load leveling 337
- management 328
- management tools 332
- massively parallel processing (MPP) 315
- message passing 316
- Moshe Bar 337
- MPP (massively parallel processing) 314
- Multicomputer Operating System for High Performance Cluster Computing (MOSIX) 336
- Myrinet 329
 - Myricom 345
- National Center for Supercomputing Applications (NCSA) 345
- networks 327
- nodes 324
- NonStop Clusters for UnixWare 337
- on demand computing 322, 349
- Open Cluster Group 337
- OpenGFS 337
- openMosix 337
- operating system 331
- parallel I/O (MOPI) 336
- Parallel Sysplex 323
- Parallel System Support Programs (PSSP) 323
- parallelism 316
- process migration 333
- protein folding 318
- pSeries 191
- radio astronomy 318
- redundancy 314
- resource sharing 316
- role of Linux clusters in emerging trends 349
- roles 316
- scalability 315
- Scyld Computing Corporation 336
- server consolidation 321
- server farm 321, 333
- Single System Image (SSI) 317, 333
- SP 323
- SteelEye Lifekeeper 345
- support software 331
- symmetric multiprocessor 323
- symmetrical multiprocessing 314
- user-mode Linux (UML) 337
- Cluster 1350 144, 324, 338
- cluster interconnect 328
- Cluster Systems Management (CSM) 332
- Clustered Enterprise Servers (CES) 324
- ClusterProven 342
- Code Versioning System (CVS) 55
- collaboration 21
- commodity off-the-shelf (COTS) 335
- common network technology 329
- Common Public License (CPL) 57
- communication types 328
- competitive software
 - data management 284
 - WebSphere 276
- compiler 47
- compute node 325, 327
- Conectiva 34
- Conectiva Enterprise Server 8 210
- consolidation 138, 190
 - iSeries 199, 211
- control node 325
- COTS (commodity off-the-shelf) 335
- coupling 318
- CSM (Cluster Systems Management) 332
- customer references for IBM Software 301
 - Diveo Broadband Networks, Inc. 305
 - Grede Foundries, Inc. 305
 - Tamkang University 306
- Customer Relationship Management (CRM) 22

D

- data management
 - competitive software 284
 - services for Linux 367
- data storage 19
- database 20
- DB2
 - for Linux Migration Services 367
 - support for Linux 367
- DB2 Product Family 279
- DB2 UDB support for Linux 281
- DB2 Universal Database 149
- DB2 Universal Database Enterprise-Extended Edition (EEE) 334
- Debian Linux 42
- DHCP 19
- direct I/O 203
- distributed computing environment 318
- distributed enterprise 214

distribution 4, 24
iSeries Linux 209

distro 24

DNS (domain name server) 19

domain name server (DNS) 19

dynamic LPAR (DLPAR) 181

E

E*TRADE Financial 147

Eclipse 58

ECM (enterprise content management) 21

Edition 234

e-mail antivirus scanners 9

emerging cluster roles 321

enterprise application 153

enterprise content management (ECM) 21

Enterprise Edition 234

Enterprise Resource Planning (ERP) 22

Enterprise X-Architecture 129

eOne Group 23

eOneCommerce 23

F

failover 320

Fast Ethernet 329

FAStT200 330

Fibre Channel 20, 330

Fibre Connection (FICON®) 323

file and print serving 19

file structure 15

file system 331

firewall 19

Free Software Foundation (FSF) 49

Free Standards Group 30, 57

Linux Standard Base 57

FSF (Free Software Foundation) 49

G

GB18030 standards 33

gcc 46

General Parallel File System (GPFS) 332

General Public License (GPL) 30, 33, 50

Gigabit Ethernet 329

GNOME 28, 57

GNU 54

GNU C Compiler (gcc) 16, 54

GNU General Public License 30

GPFS (General Parallel File System) 332

GPL (General Public License) 30, 33

grid computing 346

H

HAC (high-availability computing) 319

hardware abstraction layer 16

hardware enablement for pSeries 184

Hardware Management Console (HMC) 182

head node 327

heartbeat program 335

high availability 18

high availability cluster 335

High Performance Computing (HPC) 316, 334
application examples 316

High Throughput Computing (HTC) 318, 332

high-availability cluster 142

high-availability computing (HAC) 319

high-availability solutions 157

high-performance clusters services 171

high-performance computing 144

HMC (Hardware Management Console) 182

host 203

Host Bus Adapters (HBAs) 330

HPC (High Performance Computing) 316, 334

HPC appliance 349

HTC (High Throughput Computing) 318, 332

hypervisor 182

I

I/O device, adapter support and availability 184

IBM

Global Services 352

open source 56

open source licenses 57

open source software 57

services for Linux 351

Services introduction 352

IBM Director 138

IBM JDK Linux 193

IBM Linux Integration Centers 309

IBM Public License (IPL) 57

IBM Software

DB2 279

for Linux 266

Lotus Domino 292

Tivoli 286

IBM Support Line 356

- IBM TotalStorage Enterprise Storage Server 330
- IDS (Intrusion Detection Services) 19
- IFS (integrated file system) 198
- InfiniBlue 4X Host Channel Adapter (HCA) 331
- Informix product family 282
- infrastructure 18
 - applications for iSeries 101
- infrastructure solution 161
- Inner Circle 56
- install node 326
- Integrated Development Environment (IDE) 58
- integrated file system (IFS) 198
- Integrated xSeries Adapter 212
- Integrated xSeries Server 212
- integration
 - iSeries 200
 - iSeries Linux 204
- Intel Itanium processor 345
- interprocess communication 328
- Intrusion Detection Services (IDS) 19
- iSCSI 20, 330
- iSeries 197
 - application solutions 218
 - consolidation 211
 - I/O 203
 - infrastructure applications 101
 - Integrated xSeries Adapter 212
 - Integrated xSeries Server 212
 - iSeries Linux information matrix 201, 370
 - iSeries services for Linux 370
 - Linux integration 204
 - Linux services for the iSeries 370
 - models 204
 - performance 200
 - Portable Application Solutions Environment (PASE) 212
 - reliability 200
 - security 201
 - service solutions 234
 - Software Group products 215
 - solutions 230
 - total cost of ownership 199
- iSeries distributions
 - Red Hat 210
 - SuSE 210
 - Turbolinux 210
- iSeries Editions 234
- iSeries Technology Center 234
- iSeries Test Drive for Linux 100, 218

- ISV and application solutions 150
- Itanium 345
- iTC 234

J

- J.D. Edwards 23
- jam session 8
- Java 46
 - iSeries 230
- JDK Linux 193

K

- K Desktop Environment (KDE) 28, 57
- KD (knowledge discovery) 21
- KDE (K Desktop Environment) 28
- kernel 4, 15, 43, 55
- KM (knowledge management) 21
- knowledge discovery (KD) 21
- knowledge management (KM) 21

L

- LCS (Linux Virtual Server) 336
- legacy applications 366
- Li18nux standard 33
- license
 - closed 49
 - free software 50
 - shareware 49
 - software 48
 - X 51
- line-of-business applications 22
 - iSeries 101
- Linus Torvalds 43, 54
- Linux
 - AIX 5L 177
 - application solutions 149
 - commodity off-the-shelf (COTS) 335
 - execution on PowerPC 188
 - high availability 57, 335
 - how it runs on the pSeries server 174, 238
 - I/O flexibility 203
 - IBM Software 266
 - IBM Software on the pSeries server 192
 - Integrated Feature for 366
 - iSeries integration 204
 - kernel source 188
 - load balancing 336

- on the pSeries 670 181
- on the pSeries 690 181
- on the pSeries LPAR 181
- operating system 14
- pSeries sources of information 187
- RAS capabilities 186
- relationship with AIX 184
- role of emerging clusters 349
- scalability on the pSeries 185
- Web site platform matrix 372
- WebSphere 271
- Linux cluster 142
- Linux distribution 24
- Linux for S/390 (31-bit) 70
- Linux on xSeries IBM Software 149
- Linux on zSeries (64-bit) 70
- Linux Solution for e-business 170
- Linux Standards Base 30
- Linux Technology Center (LTC) 57
- Linux Virtual Server (LVS) 336
- Linux, UNIX, and Windows 59
 - total cost of ownership 60
- LinuxBIOS Project 17
- Linux-capable pSeries model structure 174
- Linux-related WebSphere resources 276
- LinuxWorld Conference and Expo 338
- load balancing 321, 333, 336
- load leveling 337
- load-balancing cluster 143
- Lotus
 - IBM Linux services for 369
- LPAR 181–182
 - iSeries 203
- LTC (Linux Technology Center) 57

M

- management node 326
- management operations 328
- management tools 332
- manager jam 8
- massively parallel processing (MPP) 314–315
- metacluster 346
- middleware 20
- Minix 3
- MIT 51
- monitoring 9
- MOPI 336
- MOSIX (Multicomputer Operating System for High

- Performance Cluster Computing) 336
- MPP (massively parallel processing) 314–315
- Multicomputer Operating System for High Performance Cluster Computing (MOSIX) 336
- Myricom 192
- Myrinet 192, 329

N

- NASA 334
- National Center for Supercomputing Applications (NCSA) 345
- National Digital Mammography Archive (NDMA) 346
- NCSA (National Center for Supercomputing Applications) 345
- NDMA (National Digital Mammography Archive) 346
- Netscape 51
- networks 327
- node 324
 - compute 325, 327
 - control 325
 - function 324
 - head 327
 - install 326
 - management 326
 - storage 325, 327
 - type 327
 - user 326
- non-IBM middleware 152
- NonStop Clusters for UnixWare 337

O

- object-code-only (OCO) 45
- OCO (object-code-only) 45
- on demand computing 322, 349
- Open C compiler, gcc 46
- Open Cluster Group 337
- open source 4, 29, 49
 - IBM 56
 - IBM licenses 57
 - IBM software 57
 - initiative 51
 - software 46
- Open Source Cluster Application Resources (OS-CAR) 57, 212, 337
- Open Source Definition (OSD) 51
- Open Source Developer Lab (OSDL) 57

- Open Source Initiative (OSI) 57
- OpenLDAP 57
- openMosix 337
- OpenOffice 229
- OS/400 198
 - DB2 UDB 198
 - IFS 198
 - V5R1 198, 201–202
 - V5R2 202
- OSCAR (Open Source Cluster Application Resources) 212, 337

P

- parallel computing 316
- parallel I/O 336
- Parallel Sysplex 323
- Parallel System Support Programs (PSSP) 323
- parallelism 316
- partitioning 202
 - logical 198
- Pay as You Grow 139
- PCI hot plug 57
- PDA (personal digital assistant) 24
- percolation 53
- performance monitoring 10
- performance of iSeries 200
- personal digital assistant (PDA) 24
- personal Linux education plan 369
- plug-in 58
- PowerPC, iSeries 209, 212, 218
- PPC-32 57
- PPC-64 57
- process migration 333
- protein folding 318
- pSeries
 - consolidation 190
 - hardware 184
 - introduction 173, 237, 377
 - Linux application solutions 192
 - Linux clusters 191
 - Linux for pSeries offerings 371
 - Linux running on 174, 176
 - operating system scenarios 176
 - running Linux 174, 238
 - scalability 185
 - software for Linux 192
- PSSP (Parallel System Support Programs) 323

R

- RAID 20
- RAS (Reliability Availability Serviceability) 186
- Red Flag Linux 41
- Red Flag Software 41
- Red Hat 30
- Red Hat - IBM announce 136
- Red Hat Enterprise Linux 3 (RHEL3) 136
- Red Hat Package Manager (RPM) 30
- Redbooks Web site 387
 - Contact us xxiii
- redundancy 314
- reliability, iSeries 200
- remote I/O 139
- Richard M. Stallman (RMS) 46
- router 19
- RPM (Red Hat Package Manager) 30
- RS/6000 323

S

- Samba 57, 161
- SAP 23
- scalability 315
- scalability on the pSeries 185
- Scyld Computing Corporation 336
- security 154
 - iSeries 201
- server consolidation 18, 138, 190, 211, 321
- server consolidation service 171
- server farm 333
- service processor 326
- services
 - and IBM Software 365
 - IBM Linux Learning services 369
- shareware software 49
- shell 15
- Single System Image (SSI) 317, 333
- Single Systems Image project 337
- SLES (SuSE Linux Enterprise Server) 36
- SMP (symmetric multiprocessing) 314
- software for Linux 192
- Software Group products 215
- software installation 328
- source code 46
- SP 323
- SSI (Single System Image) 317, 333
- static LPAR 181
- SteelEye LifeKeeper 345

- storage access 329
- storage node 325, 327
- Stream Control Transmission Protocol 57
- Supply Chain Management (SCM) 22
- support
 - IBM Linux Support Line 369
- Support Line 356
- SuSE 35
- SuSE Linux Enterprise Server (SLES) 36
- symmetric multiprocessing (SMP) 314
- symmetric multiprocessor 323
- systems management 21

T

- TCO (total cost of ownership) 60, 199
- TCOnow! 199
- Telecom Application Server 334
- thick client 23
- thin client 23
- Tivoli 21, 149
 - Linux supported Tivoli products 368
 - on Linux 286
 - Tivoli for Linux 368
- Tivoli for Linux services 368
- total cost of ownership (TCO) 60, 199
 - findings 60, 63
- TotalStorage Enterprise Storage Server 330
- TotalStorage IP Storage 200i 330
- TPC-H 280

U

- U2 product family 283
- UML (user-mode Linux) 337
- UnitedLinux 32
- Universal Database 20
- UNIX 21, 49, 54, 336
- Update Xpress 138
- USB 57
- USENIX 57
- user node 326
- user-mode Linux (UML) 337

V

- virtual I/O 203
- Virtual or physical partitioning 139
- VMware 139
- VMware ESX Server 140

- VSWITCH 86

W

- Web serving 19
- Web-enablement applications for iSeries 102
- WebSphere 21, 271
 - competitive software 276
 - product services matrix for Linux 366
 - resources for Linux 276
 - services for Linux 366
 - software platform 149
- WebSphere Application Server 23, 334
- WebSphere Commerce Suite 334
- WebSphere Developer Domain 277
- WebSphere Edge Server 334
- WebSphere Everyplace Suite 334
- Why Linux 199
 - iSeries 199
- Windows 23
- workload consolidation 138

X

- X-Architecture 128
- X-Architecture technology 133
- xCat 332
- xSeries 127
 - architected solutions 166
 - services solutions 170
 - xSeries Linux information matrix 133, 371
 - xSeries services for Linux 371
- xSeries 300 332

Y

- Yellow Dog Linux 41

Z

- zSeries
 - zSeries Linux information matrix 74, 370
 - zSeries services for Linux 370



Redbooks

Linux Handbook

A guide to IBM Linux Solutions

(1.0" spine)
0.875" <-> 1.498"
460 <-> 788 pages



Redbooks

Linux Handbook

A guide to IBM Linux Solutions and Resources

Review the latest IBM Solutions for Linux

Understand the Linux implementation for each IBM eServer platform

Review the extensive resources available from IBM for Linux solution implementation

IBM® has a strong commitment to Linux as an operating system as well as to Linux-based solutions. This IBM Redbook gives a broad understanding of IBM products and how they relate to the Linux operating system. It tries to encapsulate the IBM Software Solutions that are available for Linux-based servers. It also describes the Linux enablement throughout the IBM *eServer* product line.

This second edition includes a new chapters on Linux Solutions and IBM Storage in a Linux environment. All the other chapters have been revised to cover new IBM announcements, and the amazing growth in Linux adoption.

This IBM Redbook offers enough information to give you a high-level understanding of Linux. Therefore, it is not meant to be a highly technical document. It also provides pointers to further sources of information.

This redbook is intended for IBM representatives, business partners, and customers who are involved in considering and planning Linux-based software solutions. You will find this redbook particularly helpful if you are planning to test or implement Linux in an enterprise.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks